

Statement of Need (SON)

C.1 Procurement Objective

The objective of this procurement is to provide the computational and associated resources necessary to support continued advances in environmental modeling capabilities and other high-performance computing system requirements that may arise within NOAA and at other partner agencies. This procurement will be known as the NOAA Research and Development High Performance Computer System (R&D HPCS) acquisition.

This procurement embodies a new approach that NOAA adopted to holistically manage its HPCS as a corporate asset. Prior to this procurement, NOAA organizations requiring HPC procured them along organizational lines. NOAA determined that its HPCS will be established and managed based on functional requirements. NOAA established two functional requirements to meet its HPC needs: (1) Operations and (2) Applied Research and Development. This Statement of Need describes the functional requirements of NOAA's Applied Research and Development.

A primary goal of the new approach is to achieve economies of scale in conducting the R&D HPCS acquisition while providing maximum flexibility in the resulting contract document. NOAA intends to achieve economies of scale primarily through consolidation of requirements into fewer acquisitions. In order to achieve maximum flexibility, NOAA's intention for its R&D HPCS contract is to be broad enough in scope and to have sufficient reserve capability (contract ceiling) to accommodate both existing and unanticipated requirements.

The contract supporting this acquisition consists of a four-year base period, a four-year option period, and a one-year option to provide for contract transition (see Section C.9.1 for more information about the option periods). The Government expects that initial delivery, for Subsystems associated with workstreams 7-9, will occur in October 2005. The Government expects delivery, for any additional Subsystems or Components associated with workstreams 1-6, will occur in October 2006. A workstream (WS) is a single instance of end-to-end processing (including pre- and post-processing).

The Government requires a single Contractor to be responsible for the design, installation, maintenance, and support of the HPCS. The HPCS offered shall meet the stated objectives and specifications set forth in this SON and shall include all hardware and software necessary to operate as a complete, functional, balanced, and highly reliable system. The HPCS is comprised by Subsystems that are designed to help accomplish NOAA's R&D missions, as represented by workstreams. Some of the components that comprise a Subsystem include: Large-Scale Computing, Development Computing, Post-Processing and Analysis Computing, Data Management, Hierarchical Storage Management, Interconnects, Software, and Visualization. A single Contractor shall serve as the point-of-contact for the entire HPCS, even though the HPCS may involve components from a number of different vendors, and potentially be located at

multiple sites. Fundamentally, the Government must improve all aspects of NOAA's computing environment in order to fulfill its R&D mission. These aspects include:

- High-performance computing, including large-scale computing, large-scale data post-processing, analysis, and visualization capabilities. As described in this document and Section J, the R&D workstreams are comprised of computationally intensive environmental modeling applications coupled to I/O-intensive codes and extensive data storage.
- Hierarchical Storage Management System (HSMS). The HSMS offered shall provide archiving capacity to meet the expected rates of data production on the HPCS.
- Software for resource management, system administration, and application development. The Government requires operating systems and cluster software that can manage resources. Complete and functional FORTRAN90, C, and C++ application development environments shall also be provided.
- Reliability, availability, and support. Availability of at least 96% (24 hours/day, 7 days/week, calculated each month) has been the historical goal for NOAA's R&D HPCS. The HPCS offered must continue NOAA's historically high utilization of its computing resources. System reliability, availability, and Contractor support are considered fundamental aspects of the HPCS.

Of the annual funding, 94% will be dedicated to the components and services for the HPCS. Under direction from the Government, the remaining 6% of the funds will be reserved to refine key areas of the HPCS or other aspects of the computing environment, including facilities related items, covered under the scope of the contract. Improvements will be made for performance, efficiency, [security](#) or usability of the overall system. These areas may include, but will not be limited to, node, disk, memory, visualization, server, the network infrastructure, and additional support. The amount of the reserve funds may be unilaterally increased by the Government at any time during the course of the contract. The Government is not under obligation to direct the 6% reserve to the Contractor. Key areas will be identified on an as-needed basis by performance assessments, including an annual system performance review by the Government. The Government and the HPCS Contractor will work together to identify the necessary items that will best meet NOAA's computing needs. Actual purchases for this purpose will be solely at the Government's discretion.

The capabilities of the large-scale computers, hierarchical storage management system, analysis and visualization platforms, and network bandwidth shall be well-matched in a way that minimizes bottlenecks to the flow of information while maximizing performance. Achieving this proportionality in the acquired capabilities is an essential goal of this procurement. The Contractor shall provide balanced performance.

The computing resources available to NOAA must meet its scientific needs throughout the life of the contract, so the Government requires a phased delivery of all components of the HPCS. The initial delivery of the HPCS shall provide an increase over current capabilities in computational throughput for NOAA R&D. At least one significant

upgrade to the sustained throughput shall be provided during the base contract period, with archiving and other HPCS capabilities increasing commensurately. Individual components of the HPCS need not be upgraded simultaneously.

The provisions of this SON establish the Government's minimum acceptable capabilities of the HPCS based on NOAA's experience in performing its mission. However, innovation in proposed high-performance solutions is encouraged. Newer technologies or an approach different from that presented here may provide opportunities to further increase performance or enhance efficiency.

C.2 Background and Purpose

The mission of the National Oceanic and Atmospheric Administration (NOAA) is to understand and predict changes in the Earth's environment and to conserve and manage coastal and marine resources to meet our Nations' economic, social, and environmental needs. NOAA's mission is embodied in its four strategic goals:

- Protect, restore, and manage the use of coastal and ocean resources through ecosystem-based management.
- Understand climate variability and change to enhance society's ability to plan and respond.
- Serve society's needs for weather and water information.
- Support the Nation's commerce with information for safe, efficient, and environmentally sound transportation.

In support of NOAA's mission and goals, NOAA conducts research and gathers data about the global oceans, atmosphere, space, and sun, and applies this knowledge to science and services that touch the lives of all Americans.

The ability to "predict changes in the Earth's environment" depends primarily on a diverse set of environmental models, requiring considerable computational resources. These models are developed through research and development (R&D) efforts within NOAA that occur primarily at three organizations: Office of Oceanic and Atmospheric Research's (OAR) Forecast Systems Laboratory (FSL), located in Boulder, Colorado; National Weather Service's (NWS) National Centers for Environmental Predictions (NCEP), located in Camp Springs, Maryland; and OAR's Geophysical Fluid Dynamics Laboratory (GFDL), located in Princeton, New Jersey. Computational modeling continues to emerge throughout NOAA and is germane to the daily operation of many of its laboratories. Organizations such as the Pacific Marine Environmental Laboratory (PMEL), Climate Diagnostics Center (CDC), Aeronomy Laboratory (AL) and Air Resources Laboratory (ARL) use models for study of atmosphere, ocean, climate, air quality and ecosystem behavior.

C.3 Current NOAA R&D HPCS

The current HPCS is comprised of three Subsystems with components that include an IBM Power 4 cluster; a Linux cluster of Intel Xeon single and multi- core processors interconnected by Myrinet; and SGI Origin 3800, Origin 3900 and Altix systems. The

entire HPCS includes over 7.5 PB of storage capacity and greater than 15 million files. See Appendix A for a more complete description of the current HPCS.

C.4 Benchmarks

C.4.1 Workstream (WS) Profile

As previously mentioned, a workstream (WS) is a single instance of end-to-end processing (including pre- and post-processing). The set of instances of a given workstream within the Throughput Benchmark (see Section J) represents a set of jobs with similar characteristics within NOAA's current HPC infrastructure. NOAA's R&D HPCS Project Team defined 9 representative workstreams. See Section J for more complete descriptions of representative workstreams.

WS1	CM2-ESM	Coupled Earth System Model
WS2	CM2-HR	High Resolution Coupled Climate Model
WS3	HIMF	Very High Resolution Ocean Model
WS4	EMTB	Environmental Modeling Test Bed
WS5	CMDC	Climate Model Development and Calibration
WS6	DAD	Data Assimilation Development
WS7	RUC	Rapid Update Cycle
WS8	WRF-CHEM	Weather Research and Forecast w/ atmospheric chemistry
WS9	WRF-SI	Weather Research and Forecast w/ static initialization

C.4.2 Benchmark Performance

Within a workstream's funding profile (see Section C.4.3), the Government requires the maximum System Life Throughput (SLT) utilizing a balanced solution for each workstream in the contract period. The Government considers every workstream to be of equal importance and expects significant performance increases for all of them.

The benchmark is comprised of a set of application workstreams. Each workstream is a sequence of steps designed to represent the complete end-to-end execution of a single modeling application used in NOAA R&D. Each workstream may contain multiple applications with a mix of compute-bound and I/O-bound codes. The Contractor shall minimize the execution time of each workstream and maximize the overall throughput when multiple workstreams are run concurrently, as described in Section J.

SLT is the metric of system performance for each workstream. The workstream SLT is a measure of system performance and availability integrated over time. See Section C.6.1.2 for more on workstream SLT.

The workstreams were baselined on three different computer systems. WS1 - WS3 were baselined on NOAA's Origin 3000 HPC in Princeton, NJ; WS4 - WS6 were baselined on NOAA's IBM Power 4 Regatta H HPC in Washington, D.C.; and WS7 -

WS9 were baselined on NOAA's Xeon-based HPC cluster in Boulder, CO. Further details of these workstream baselines can be found in Section J.

The Government can most effectively use performance increases that are implemented at regular, but not too frequent intervals during the contract. Further, the Government would have difficulty using disproportionately high performance delivered either very early or late in the contract period. Accordingly, a given workstream shall receive no more than two (2) performance upgrades during either the base or option periods.

The Government will evaluate performance increases in terms of the performance baseline described in Section J. The initial deliveries for WS7 - WS9 are required in Q1FY2006, the initial deliveries for WS1 - WS6 are required in Q1FY2007. The Government requires maximum System Life Throughput obtained by at least one significant mid-life performance upgrade during both the base and option periods. *It is desirable that the mid-life performance upgrades for the base period for WS7 – WS9 occur during Q4FY2007 or Q1FY2008. It is also desirable that the WS1 – WS6 performance mid-life upgrades should occur during Q2 or Q3FY2008.*

The Government recognizes the value and benefits of optimized application software and improved compilers, libraries and runtime environments. On the other hand, the primary goal of the mid-life and option period upgrades is to produce significant workstream performance enhancement through the addition of substantial subsystem hardware. *Further, benchmarks may change by mutual agreement between Government and Contractor at mid-life and option period upgrades.* Therefore, benchmarking for mid-life and option period upgrades shall be handled as follows:

- *For mid-life and option period upgrades where benchmark revisions are not defined:*
 - o *Changes to the benchmark software for performance relative to the previous Live Test Demonstration (LTD) will be allowed at the sole discretion of the Government.*
 - o *Changes in compiler and/or any other runtime environment and/or library with respect to options and/or versions relative to the previous LTD will be allowed at the sole discretion of the Government.*
 - o *Any changes that are allowed will be used to create a new baseline with respect to the existing system.*
- *For mid-life and option period upgrades where benchmark revisions are defined:*

- o Revised workstream benchmark codes and performance baseline data will be delivered by the Government to the Contractor no later than 6 months prior to delivery of an upgrade. The Contractor may then perform code, compiler and runtime environment and library optimizations and include those changes from the baseline in the proposed system upgrade and SLT provided those optimizations are acceptable to the Government.
 - o A performance increment is always defined with respect to a set of source code and a well defined set of compiler, runtime environment and library options and versions acceptable to the Government..
- Any changes to the benchmark as outlined in this paragraph must be described in detail and in writing to the Government one month prior to the delivery of the upgrade.
- Revised workstream benchmark codes and performance baseline data will be delivered by the Government to the Contractor no later than 6 months prior to delivery of an upgrade.
- Throughput and Scaling performance data shall be reported as described in Section J:
 - o Throughput and Scaling values with class A changes only.
 - o Any additional Throughput and Scaling values with class B, C and D changes. All such changes will be evaluated for risk and approved in the same manner as the RFP response.
- The configuration management (CM) process will be utilized as the forum to discuss scientific advancement as it relates to model performance as well as software architecture changes to accommodate the new science. In the context of the site-specific CM meetings, this framework will keep both Government and Contractor personnel up to date with desired future trends and aid the benchmark revision process. Among other items, the CM forum will be used by the Contractor to suggest benchmark code changes to NOAA technical staff for technical evaluation.

C.4.3 Funding Profile

Table I shows the projected funding levels associated with each workstream. These levels indicate current estimates of NOAA support for the programs represented by the workstreams. Contractors are expected to use this information in establishing performance goals for the benchmark. The offered configuration shall provide flexibility to the Government to meet changing priorities throughout the contract period. The funding begins for workstreams 7-9 in FY2006. The funding for the remaining workstreams will begin in FY2007. The Government will retain 6% of the

cumulative annual “Total” funding amount (the far right-hand column) cited in Table I below as a reserve. Based on information provided by the Contractor, the Government will evaluate proposals to verify that workstream performance is appropriate to the funding profile.

The available funding level for the one-year extension options is anticipated to be approximately half of the normal yearly expenditures. More information about the extension options can be found in Sections C.9.2 and C.9.3.

Table I - Funding Profile in \$Millions (shown before 6% reserve is removed)

		WS1	WS2	WS3	WS4	WS5	WS6	WS7	WS8	WS9	Total
Base	FY2006							\$1.3	\$1.0	\$0.7	\$ 3.0
	FY2007	\$4.0	\$6.0	\$4.0	\$1.933	\$1.933	\$1.934	\$1.3	\$1.0	\$0.7	\$ 22.8
	FY2008	\$4.0	\$6.0	\$4.0	\$1.933	\$1.933	\$1.934	\$1.3	\$1.0	\$0.7	\$ 22.8
	FY2009	\$4.0	\$6.0	\$4.0	\$1.933	\$1.933	\$1.934	\$1.3	\$1.0	\$0.7	\$ 22.8
Option	FY2010	\$4.0	\$6.0	\$4.0	\$1.933	\$1.933	\$1.934	\$1.3	\$1.0	\$0.7	\$ 22.8
	FY2011	\$4.0	\$6.0	\$4.0	\$1.933	\$1.933	\$1.934	\$1.3	\$1.0	\$0.7	\$ 22.8
	FY2012	\$4.0	\$6.0	\$4.0	\$1.933	\$1.933	\$1.934	\$1.3	\$1.0	\$0.7	\$ 22.8
	FY2013	\$4.0	\$6.0	\$4.0	\$1.933	\$1.933	\$1.934	\$1.3	\$1.0	\$0.7	\$ 22.8
Total		\$28.0	\$42.0	\$28.0	\$13.531	\$13.531	\$13.538	\$10.4	\$ 8.0	\$ 5.6	\$162.6

C.5 HPC Subsystem Components

C.5.1 Computing Requirements

C.5.1.1 Large-Scale Computing (LSC) Component

The Contractor shall provide a Large-Scale Computing (LSC) component at a substantial increase in sustained throughput over NOAA’s current supercomputers described in Appendix A. Sustained throughput shall be measured by a throughput benchmark (see Section J) comprised of workstreams that are surrogates for NOAA's expected future workload. The metric of performance for the LSC is based on sustained throughput. The scalability of the computational platform(s) shall be measured by a benchmark designed to reveal the performance and scaling characteristics of individual codes as they are executed on different processor core counts. (Scalability is with respect to different socket counts where applicable). The scalability measure will assist the Government in evaluating performance projections.

It is required that whenever any set of resources (such as a cluster node) in the LSC fails, batch jobs using those resources must be capable of being rerun without user intervention. In this situation, only interactive sessions hosted on the failed resources will be lost, and the Subsystem shall allow users to continue to be able to login interactively. It is desirable that failover be to processor architectures that are binary-compatible with and running the same operating system and level as the failed resources. The capability of the LSC to operate in degraded mode during repairs is required. It is desirable that the LSC have no single point of failure. The Government requires an availability level (defined in C.6.1.2) of at least 96% on every Subsystem in the R&D HPCS.

At least one substantial upgrade to the sustained throughput of the LSC, as measured by a throughput benchmark, shall be provided during the base contract period. This has traditionally been considered the mid-life upgrade.

Prior to the initial delivery, the Government requires an early access system identical in architecture and software environment to that offered in support of workstreams 7-9. The system need not be as large as the proposed system. The early access system should be available within 21 days of contract award. This system must be dedicated to the use of the Government until acceptance testing begins. The Government must be allowed privileged access to the early access system. The Government expects to begin configuring the system including, but not limited to, creating accounts and configuring the offered batch system. The Government will also port and test critical applications.

C.5.1.2 Interactive Component

For each LSC Subsystem, binary compatible interactive development resources shall be provided. The interactive environment shall have the same software environment as the LSC Subsystem. The interactive resource shall represent 5% to 15% of the LSC component. It is desirable that the amount of interactive resource available is adjustable by the Resource Management Software (see Section C.5.3.1) without a system reboot.

The Government desires that the operation of interactive resources have a minimal impact on the batch production resources and vice-versa. It has been NOAA's experience that the nature of interactive R&D work creates resource contention with batch production jobs, and batch production jobs slow interactive response time. The Government desires the ability to reassign batch production resources to interactive work, and vice versa, without a reboot or a restart to the batch system.

The Government desires early access to the interactive resources for training, software porting and tuning, and OS and application configuration testing. Early access to initial system delivery and subsequent system upgrades (with some Contractor provided engineering assistance) up to 6 to 9 months in advance of delivery is desirable.

C.5.1.3 Post-Processing and Analysis Component

Some of the HPC Subsystems described in Appendix A execute compute-bound and I/O-bound codes separately. For example, the climate computing resources, located in Princeton, use a Large Scale Cluster for the former and an Analysis Cluster, configured for very large sustained I/O, for the latter. The Contractor shall provide resources that can efficiently execute both types of codes, with the goal of minimizing the execution time of each workstream and maximizing the overall throughput when the workstreams targeted for that subsystem are run concurrently.

C.5.2 Data Management Requirements

The Government's scientific data are critical assets. As noted in section C.6, the Government requires the highest level of data integrity and at least 99% availability for access to its scientific data. The Government requires this access to its scientific data in the absence of the Large-Scale Computing resources.

C.5.2.1 Home File System (HFS)

A Home File System (HFS) is required for each HPC Subsystem provided. This file system will have small quotas and be used to store source-code, small data sets, and environment initialization files. The Government desires initial HFS storage of several GBs per user with the ability to grow with each upgrade. The HFS shall be globally visible to the specific Subsystem with which it is associated. If the HFS is not available, no workload can be accomplished by the associated Subsystems. If no workload can be accomplished, the Subsystem will be in downtime, as described in Section C.6.1. The HFS shall be backed up utilizing an automated method (see Section C.5.2.8).

C.5.2.2 Fast-Scratch File System

A Fast-Scratch file system is required and shall provide high-bandwidth local storage that is visible by all sockets within a given batch job. The Fast-Scratch file system will be purged frequently. The Fast-Scratch file system will not be backed up. Initially, per Subsystem, the Fast-Scratch file system shall be able to support 100 million files and individual files that are up to 500 GB in size. The Government expects that these requirements will grow over the life of the contract.

C.5.2.3 Long-term Scratch File System

Long-term disk storage is required and shall provide a staging area for users to temporarily place files that will continuously be manipulated. Large data sets that are not packed into a larger file are not conducive to a HSMS file system. This file system does not need to be backed up and can be purged of aged data. Initially, per Subsystem, the Long-term Scratch file system shall be able to support 100 million files and individual files that are up to 500 GB in size. The Government expects that these requirements will grow over the life of the contract.

It is anticipated that the scientific data for WS4 – WS9 will be primarily accessible through the Long-term Scratch File System. The Government requires accessibility to its scientific data stored on the Long-term Scratch File System in the absence of the Large-Scale Computing resources.

C.5.2.4 Hierarchical Storage Management System

The Hierarchical Storage Management System (HSMS) component will be purchased. The Government desires a flexible HSMS hardware configuration that can be changed or enhanced to optimize HSMS performance. Upgrades to the HSMS shall be provided commensurate with upgrades to the computational resources, with overall system balance the goal. The HSMS shall meet the

performance requirements of the HSMS Archive benchmark given in Section J.1.4.4.

The Government requires accessibility to its scientific data stored on the HSMS in the absence of the Large-Scale Computing resources. It is anticipated that the scientific data for WS1-WS3 will be primarily accessed through the HSMS.

The Government requires a minimum of a two-tiered storage scheme for its data archive, comprised of frequently accessed nearline storage (disk storage or tape storage robotically mounted at high speed) and offline storage (tape storage robotically mounted). This will effectively satisfy the requests for scientific data that make up NOAA's R&D workload (as discussed in section C.4). The HSMS software shall provide automatic migration between data archive tiers based on a combination of access time and file size. The Government desires that files recalled from the offline tier automatically revert to nearline tier residency. User-specified migration between tiers is also desirable. The Government desires a file display command, comparable to UNIX "ls -l", which shows each file's residency in the nearline or offline tier. The Government requires a daily report, submitted to the COTR, of storage used by groups and users in all storage tiers.

If disk storage is used for the nearline tier it shall be fault-tolerant and sized for the Government's file size distribution and data access patterns. As it may be difficult to recover from disk storage failure, the Government requires all HSMS disk files to either have a secondary copy or a backup within 24 hours.

The Government requires uninterrupted access to its legacy data (see Table IIIa). NOAA's existing data archive for workstreams 1-3 resides on a Government-owned, SGI/StorageTek HSMS in Princeton (see Appendix A). This legacy HSMS is available to the Contractor as unrestricted GFE beginning FY2007 (see Appendix C). The Government desires a transparent migration of this legacy archive to the new HSMS, maintaining the present file system image. The Contractor shall make arrangements to keep this data available during the transition, potentially at a reduced level of performance.

NOAA's existing data archive for workstreams 4-6 resides on Government-owned StorageTek Powderhorn silos located at IBM's facility in Gaithersburg, MD (2 silos) and the Government facility in Fairmont, West Virginia (1 silo) (see Appendix A). These silos presently support NOAA's Operational Central Computer System (OCCS) and its backup. These legacy HSMS silos are available to the Contractor as unrestricted GFE beginning in FY2007 (see Appendix C). The Contractor shall make arrangements to keep this data available. The Government desires a transparent migration of this legacy archive data to the new HSMS, maintaining the present filesystem image. The HSMS proposed to support workstreams 4-6 is required to interface with either the OCCS or its backup which are located at Fairmont, WV. and Gaithersburg, MD. (See C.5.2.5)

NOAA's existing data archive for workstreams 7-9 resides on a Government owned, Sun/ADIC HSMS in Boulder (see Appendix A). This legacy HSMS will remain under maintenance until September 9, 2006. At that point, the HSMS is available to the Contractor as unrestricted GFE (see Appendix C). The Government desires a transparent migration of this legacy archive to the new HSMS, maintaining the present file system image. The Contractor shall make arrangements to keep this data available during the transition, potentially at a reduced level of performance.

Multiple archives may be provided. Each data archive shall be presented to the associated subsystem as a single file system. Initially, per Subsystem, the HSMS shall be able to support 200 million files and individual files that are up to 1 TB in size. The Contractor shall provide an archive storage solution that scales in performance sufficiently to enable a transparent and cost-effective solution for at least one year beyond the expiration of the contract. The Government desires HSMS software comparable to the San Diego Supercomputer Center (SDSC) Storage Resource Broker (SRB), which transparently combines logically related small files into one larger "container" file in the archive filesystem. More information on SRB is available at <http://www.npaci.edu/DICE/SRB>. The Contractor shall provide all storage media used in the HSMS. The Government desires that files that haven't been accessed in one year will be migrated from the nearline storage to the offline storage. As directed by the Government, within 30 calendar days, the Contractor shall provide a media recovery service to recover data in the event of nearline or offline storage media failure.

For the tape storage, the Government desires a high file positioning rate, P , with matching robotic mount performance. For tape media, P (in files/hour) is given by the equation below. The "average search time" shall apply to the method the HSMS software uses to position files.

$$P = \frac{(n * 3600)}{(load + search + rewind + unload)}$$

P	=	file positioning rate
n	=	number of tape drives
$load$	=	time to load a tape in sec.
$search$	=	average search time in sec.
$rewind$	=	average rewind time in sec.
$unload$	=	time to unload a tape in sec.

For the tape storage, the Government desires several types of tape media optimized for the file size distribution, balancing file positioning time with file read time.

The Government desires that the HSMS software provide priority scheduling of data recall requests, assigning high priority to certain users or to batch jobs over interactive sessions.

C.5.2.5 Interfaces to the NOAA Operational Central Computer System (OCCS) and the Backup OCCS

The exchange of information to and from the NOAA Operations HPC is an integral part of NOAA's overall Research, Development and Operations processes; an objective of this procurement is to preserve and enhance the ability to execute such exchange of information.

The data that is generated from the OCCS, and available at both Fairmont, WV, and Gaithersburg, MD, is required to be written to the R&D HSMS that supports workstreams 4-6. This data consists of the run histories from the operational models along with observation data that are outputs from the OCCS. Table IIa shows the daily volumes that NOAA expects to be written over the life of the contract. Current practice is to "batch" these data and write them to the R&D HSMS during off-peak hours to avoid contention with other HSMS users.

The HSMS supporting workstreams 4-6 shall:

- Provide the capabilities necessary to allow for the OCCS backup system to read from and write to the HSMS.
- Provide to the OCCS backup system similar HSMS access and latency characteristics as those described in Section J.1.4.4.3

This data is composed of model run histories and data generated from pre-operational model development work.

C.5.2.6 Data Generation Profile

C.5.2.6.1 Data Generation for WS1 – WS3

NOAA's data generation for its computational baseline for WS1 – WS3 is shown below:

	WS1	WS2	WS3
Baseline Data Generation	1.4 TB/day	2.6 TB/day	1 TB/day

Data generation for WS1 – WS3 is related to the computational increases related to the workstream. Such changes shall be negotiated. Data generated (d) by the workstream (i) can be calculated using the following equation:

$$d_{i,k} = d_{i,(k-1)} \left(\frac{\overline{B}_{i,(k-1)}}{\overline{B}_{i,k}} \right)^{0.7}$$

d = Data generated by the workstream
 \overline{B} = Average workstream Benchmark time
 i = Workstream number
 kj = system configuration period

Due to changes in the computational models, this equation may need to be adjusted throughout the life of the contract. The Government reserves the right to negotiate any changes due to these adjustments.

The following example demonstrates the application of the formula above. Suppose that during the previous system configuration period ($k-1$), the workstream was generating 1 TB of data per day and the workstream Benchmark time averaged of all instances of the workstream was 10,800 seconds. Suppose further that during the current system configuration period (k), the average workstream Benchmark time is now 7,200 seconds. The data generation (d) for the current system configuration period (k) is given by:

$$d_{sample} = 1 \frac{TB}{day} \left(\frac{10800 \frac{s}{benchmark}}{7200 \frac{s}{benchmark}} \right)^{0.7} = 1.34 \frac{TB}{day}$$

C.5.2.6.2 Data Generation for WS4 – WS9

Table II shows the projected data generation associated with WS4 - WS9 throughout the contract period.

Table II -- Data Generation for WS4 – WS9 (in TB/day)

		WS4	WS5	WS6	WS7	WS8	WS9	Totals
Base	FY2006				2.4	2.2	1.6	6.2
	FY2007	2	3	2	2.4	2.2	1.6	13.2
	FY2008	4	3	4	4.8	4.4	3.2	23.4
	FY2009	5	3	5	4.8	4.4	3.2	25.4
Option	FY2010	8	6	8	9.5	8	6	45.5
	FY2011	10	6	10	9.5	8	6	49.5
	FY2012	17	6	17	12	10	8	70
	FY2013	24	12	24	12	10	8	90
Option FY2010/FY2014		24	12	24	12	10	8	90

C.5.2.6.3 Data Generation for OCCS and the Backup OCCS

Table IIa shows the projected data generation for the OCCS and the Backup OCCS.

Table IIa -- Data Generation / Access Profile for OCCS and Backup

	OCCS (Write)	Backup (Read)	Backup (Write)
FY2005	1.0 TB/day	0.8 TB/day	0.8 TB/day
FY2006	1.6 TB/day	1.2 TB/day	1.3 TB/day
FY2007	2.2 TB/day	1.7 TB/day	1.8 TB/day
FY2008	3.2 TB/day	2.5 TB/day	2.7 TB/day
FY2009	4.6 TB/day	3.5 TB/day	3.9 TB/day
FY2010	6.6 TB/day	5.1 TB/day	5.5 TB/day
FY2011	9.4 TB/day	7.2 TB/day	7.9 TB/day
FY2012	13.0 TB/day	10.0 TB/day	11.0 TB/day
FY2013	19.0 TB/day	14.0 TB/day	16.0 TB/day
FY2014	26.0 TB/day	20.0 TB/day	22.0 TB/day

Note: The totals shown in the OCCS (Write) and Backup (Write) columns of this table are in addition to the totals shown in Table II (Data Generation for WS4-WS9).

C.5.2.7 Data Retention Profile

The data life-cycle is the expected number of years that the data from a workstream will be stored. These data are all of the remaining non-scratch files from a given workstream. This includes, but is not limited to, source code, input files, run time scripts, generated output, and analysis results. The Contractor shall provide adequate storage to meet the Government's data retention requirements as specified in Tables III and IIIa.

Table III -- Data Retention Table


	Life-Cycle of the data for a given Workstream
WS1	100% as a persistent archive.
WS2	100% as a persistent archive.
WS3	100% as a persistent archive.
WS4	100% for the first year; 50% for the second year; and 15% as a persistent archive.
WS5	100% for the first year; 50% for the second year; and 15% as a persistent archive.
WS6	100% for the first year; 50% for the second year; and 15% as a persistent archive.
WS7	100% for 3 DAYS; 8% for 3 years; and 4% as a persistent archive.
WS8	100% for 3 DAYS; 8% for 3 years; and 4% as a persistent archive.
WS9	100% for 3 DAYS; 8% for 3 years; and 4% as a persistent archive.

Table IIIa -- Data Retention Table for Legacy Data

Life-Cycle of the legacy data	
GFDL	100% of legacy data as a persistent archive.
NCEP	100% of legacy data as a persistent archive.
FSL	100% of legacy data for the first year; 50% for 3 years; and 25% for 8 years.

The data retention profile for the data shown in table IIa is 100% for the first year; 50% for the second year; 30% as a persistent archive. Persistent archive is defined as the data in the archive that will exist or remain indefinitely. The funding for this storage is contained in the funding profile for workstreams 4, 5, and 6 shown in table I.

The following figure shows the accumulated data retention for WS4 – WS9. This is utilizing the data generation numbers shown in Table II of C.5.2.6 and Table III of this section. WS1 – WS3 cannot be shown without the computational growth being determined.



NOTE: Beginning ordinate of this graph is 2TB

Figure 1 - Data Retention for WS4 - WS9

C.5.2.8 Automated backup

Automated backup shall be provided for all of the following: all unique system images on the HPCS; each Home File System (HFS); and the disk-resident inodes, metadata, and files on the HSMS. Software that allows users to restore /home files via a graphical interface is desirable.

For the HFS, a combination of full and incremental backups shall be made to robotically mounted tapes. These backups shall make it possible to restore files to their state on any day during the previous two calendar months. Minimal impact of these backups on the network load is desired.

A history of bimonthly (every 2 months) full backups, for the previous 12 months, shall be produced for shelf storage. A history of annual backups shall be produced until the end of the HPCS system life for shelf storage. At the Government's discretion, shelf storage may be at a remote location. It shall be possible to restore files from these backups until the end of the file systems for which the backups are made.

All hardware and storage media used for backup shall be provided by the Contractor.

C.5.3 Software Requirements

The Government reserves the right to negotiate quantities of software. The Government further reserves the right to augment system licenses through other NOAA contract vehicles such as NOAA wide Blanket Purchase Agreements (BPAs). The Government, at its option, reserves the right to become the licensee of any software acquired under this contract on or before the termination of the contract.

C.5.3.1 Resource Management Software

Efficient operation of the HPCS requires resource management that will facilitate use of the HPCS by NOAA's scientists as well as provide maximum throughput for their workload. The Government desires to implement a variety of algorithms for monthly socket use time or enforce different resource allocations, including disk and tape quotas, for users, groups, or projects on each of the HPCS Subsystems. All Resource Management Software shall be provided with documentation, training, and an established problem resolution process. The Government requires resource management software that provides:

1. Batch queuing and scheduling (see Section C.5.3.2)
2. Operating system based accounting software, comparable to Unix SVR4 process accounting.
3. System activity monitoring software on the LSC that shows user and system CPU utilization, I/O wait time, and paging activity.
4. Disk quotas on disk space and number of files, comparable to Unix 4.3 BSD quotas.

5. System Performance Monitoring Tools to evaluate load balancing characteristics, system health, and other performance issues.

The Government desires resource management software on the LSC that provides software partitioning.

The Government desires the ability to test operating system and application software upgrades in isolation from the interactive and batch production resources on the HPCS.

C.5.3.2 Batch Queuing Software

The Government requires Batch Queuing Software to manage all HPC Subsystem computing resources. The Batch Queuing Software shall be provided with documentation, training, and an established problem resolution process.

To maintain flexibility, the Batch Queuing Software shall support execution hosts with heterogeneous socket architectures and operating systems. Heterogeneous architectures include existing GFE equipment and additional computational capabilities that may become introduced into the system that are different than the delivered system(s).

The Batch Queuing Software shall perform the following functions:

- Spool batch job scripts when submitted.
- Tag each job with a user-specified string that is at least 45 characters long.
- Run parallel-processed jobs on sockets dedicated to the job.
- Create a unique scratch filesystem temporary directory for each job.
- Enforce limits on per-process cpu time, job-total cpu time, and elapsed (wallclock) time.
- Track and display job-total cpu time as the job runs.
- Provide priority-based job scheduling, allowing higher priority to be assigned to specific jobs, or all jobs owned by a user or group.
- Limit the number of running jobs and monthly cpu hours used by users, groups, and projects or accounts.
- For accounting purposes, tag each job with a project or account string.
- Collect cluster-global job accounting information for each job, including:
 - the user, group, and project or account
 - the number of sockets dedicated to the job
 - the job-total cpu time
 - the elapsed (wallclock) time
 - the date and time of job completion
 - NOAA specified programs identifier

The Batch Queuing Software shall provide text-based displays as specified below. It shall be possible to create custom text-based displays, at a minimum by post-processing the display output with text-processing tools such as awk or perl. The Batch Queuing software shall:

- Display all cluster hosts, showing for each host
 - the host up/down status
 - the host load and resource usage
- Display all running jobs, showing for each job
 - the host(s) the job is running on
 - the number of sockets dedicated to the job
 - the job-total cpu time
 - the elapsed (wallclock) time
- Display all queued jobs, showing for each job
 - the number of sockets requested by the job
 - the date and time the job was submitted
 - the per-process cpu time limit for the job
 - the elapsed (wallclock) time limit for the job
- Display all job accounting records for an operator specified time window

The following Batch Queuing Software features are desired:

- Uniform access to all HPCS computing resources, through use of a single Batch Queuing Software package.
- Meta-scheduling capabilities able to schedule batch jobs on multiple clusters over a wide-area network.
- Priority scheduling of large jobs, keeping smaller, lower priority jobs from starting. In addition, during this draining process, scheduling of small, short jobs, which will complete before the large job starts (backfill scheduling).
- Job scheduling based on consumable resources such as host memory, host disk space, or cluster-global shared disk space.
- Job suspend, job migration, and job checkpointing features that work for FORTRAN90, C, and C++ applications using the proposed parallel programming model.
- Interactive use integrated with the Batch Queuing Software, which provides the batch execution environment to interactive sessions. This will allow batch jobs scripts to be run interactively without change. Accounting of the cpu time used by interactive sessions.
- [Provide a way for large input files to be pre-staged to the scratch filesystem before the job acquires dedicated resources, and for large output files to be post-staged after the dedicated resources are released.](#)

The following Batch Queuing Software features are useful:

- Support for the user-level commands and command options of the POSIX 1003.2d Batch Environment standard (now part of POSIX 1003.1-2001).
- Client software which allows job submission and displays directly from the Government's Linux desktop workstations without interactive login to the HPCS.
- GUI-based displays.

Failover capability for job queuing and scheduling shall be provided.

C.5.3.3 Programming Environment Software

Required application development software for each LSC, development, post-processing and analysis Subsystem of the HPCS programming environment is comprised of:

1. FORTRAN 90/95, C, C++ programming environments, including:
 - a. ANSI standard FORTRAN 90/95, C, and C++ compilers
 - b. Number of concurrent compiles per workstream group.
 - i. WS 1-3: 45
 - ii. WS 4-6: 30
 - iii. WS 7-9: 25
 - c.. macro preprocessors
 - d. source-level debuggers
 - e. performance profilers
 - f. support for 64-bit IEEE reals and integers
 - g. support for reading and writing 32-and 64-bit IEEE floating-point formats in I/O operations
 - h, MPI-1.1, MPI-2 I/O, MPI-2 one-sided communications (for Subsystems supporting parallel environments only)
 - i. The make utility
2. Facilities for source code management, including rcs and cvs.
3. OpenMP is required to the extent that processing units share memory.
4. The Etnus TotalView parallel debugger.

TotalView Configuration	WS 1-3	WS 4-6	WS 7-9
32 Process	6	7	4
128 Process			1
256 Process	4	1	

The Contractor shall describe how it will provide access to HPCS related engineering and development teams for updates and input concerning the programming environment. The Contractor, at the Government's direction, shall act as the point of contact for existing relationships with engineering and development teams. The Contractor shall attempt to establish new relationships with engineering and development teams as necessary.

Desired application software for the HPCS programming environments includes:

- 1) parallelized, optimized numerical libraries on the HPCS
- 2) optimized (possibly proprietary) I/O libraries.
- 3) data conversion facilities (for example, endian conversions and conversions of proprietary data formats used in the HPCS). Data conversion facilities shall be required if the provided solution requires heterogeneous formats.
- 4) beta versions of compilers for forward looking compiler options and extensions.

As standards evolve, the Government reserves the right to require updated versions of the software that has been approved through HPSC Configuration Management.

C.5.3.4 Commercial Off the Shelf Software

The Government holds or will hold licenses for the following COTS software. The Government requires that these applications can be installed on the Analysis Cluster (AC) and function correctly. The license server technology shall be able to support these applications. The Government reserves the right to augment the supported COTS software through this or other contract vehicles. All new requirements will go through the Change Management Process.

Software References:

- Matlab (<http://www.mathworks.com>)
- Mathematica (<http://www.mathematica.com>)
- IDL (<http://www.rsinc.com/idl/>)
- S-Plus (<http://www.insightful.com/products/splus/>)
-

C.5.3.5 Community Supported Software

Some Community Supported Software may require Subsystem specific customizations for optimal use. The Government reserves the right to require additional Community Supported Software that has been approved through HPSC Configuration Management. Required application software for the HPSC programming environments is comprised of:

- GRIB libraries and utilities (<http://www.wmo.ch/web/www/WDM/Guides/Guide-binary-2.html>)
- GNU make (<http://www.gnu.org/software/make/make.html>)
- GNU tar (<http://www.gnu.org/software/tar/>)
- Heirloom cpio (<http://heirloom.sourceforge.net>)
- libxml (<http://xmlsoft.org>)
- Environment Modules (<http://modules.sourceforge.net/>)
- netCDF and udunits libraries (<http://www.unidata.ucar.edu>)
- netCDF Operators (<http://nco.sourceforge.net>)
- Prior to the end of the Option Period, the Government expects that its modeling system software will become compliant with the Earth System Modeling Framework (ESMF). More information on ESMF is available at <http://www.esmf.ucar.edu>.
- tcsh (<http://www.tcsh.org>)

Required application software for interactive use is comprised of:

- Ferret (<http://ferret.wrc.noaa.gov/Ferret/>)
- GEMPAK/N-AWIPS (<http://my.unidata.ucar.edu/content/software/gempak/>)
- Ghostscript (<http://www.cs.wisc.edu/~ghost/>)
- Grace (<http://plasma-gate.weizmann.ac.il/Grace/>)

- GrADS (<http://grads.iges.org/grads/grads.html>)
- ImageMagick (<http://www.imagemagick.org>)
- NCAR graphics (<http://ngwww.ucar.edu>)
- Perl (<http://www.perl.com>)
- Python (<http://www.python.org>)
- Ruby (<http://www.ruby-lang.org>)
- VIS5D (<http://vis5d.sourceforge.net>)

C.5.4 Networking Requirements

The R&D HPSC requires networking services to support the following functions:

- Receipt of data from external sources
- Transfer of data to and from the R&D HSMS
- Support for user activities
- Data interface to the OCCS and backup OCCS.

Many of these functions are supported by GFE networking services (see Appendix A Tables VI, VII, and VIII). NOAA does not guarantee that the latencies over the GFE WAN connections will be adequate for interactive use. If a Contractor proposes to use any of the GFE links for interactive use, the Contractor shall analyze expected latencies to establish the adequacy of the link for the intended use. Depending on the Contractor implementation these GFE networks may have to be augmented by the Contractor to meet NOAA's requirements.

NOAA's Enterprise Network Target Architecture (ENTA) is currently being developed. The resulting NOAA-wide network infrastructure could be referred to as NOAAnet. NOAAnet may utilize cutting edge networks such as the National Lambda Rail (NLR) network. At this point in time NOAA has access rights to NLR at the Boulder facility. Upon completion of this architecture, the Government may, at its option and at any time during this contract, require the Contractor to alter the WAN connection to comply with the NOAA ENTA. Such a change will be negotiated.

OMB Memoranda M-05-22

<http://www.whitehouse.gov/omb/memoranda/fy2005/m05-22.pdf> states that agencies must use IPv6 and interface with this infrastructure. At all appropriate interfaces to the HPSC, the Contractor shall use equipment capable of utilizing IPv6 in compliance with the memorandum. The Government desires that the entire HPSC is capable of utilizing IPv6 prior to the June 2008 deadline.

The Contractor is encouraged to be innovative in their Wide Area Network solution in order to optimize total system performance. The Government desires to balance the Wide Area Network solution that is provided with the impact it will have on the interactive experience of remotely located users.

The following sub-sections will provide a description of how data is currently received and transported as well as describing the overall networking requirements.

C.5.4.1 Receipt of data from external sources

Initially, the available GFE WAN bandwidth will make the necessary data available within the Washington D.C. metro area, for workstreams 4 and 6, and to the Skaggs building in Boulder, for workstreams 7-8-9.

The Contractor shall be responsible for all WAN costs that exceed GFE (See Appendix A Tables VI and VII) regardless whether the connection is required to a Government or vendor provided facility.

Table IV- Data Ingest (TB/day)

	Required by WS4, WS6, WS7, WS8, and WS9	Required by WS7, WS8, and WS9 in addition to that required by the 5 workstreams in column 1	Required by WS4 and WS6 in addition to that required by the 5 workstreams in column 1
FY2006	0.75	0.25	
FY2007	1.5	0.5	1
FY2008	1.5	0.5	1
FY2009	3	1	2
FY2010	3	1	2
FY2011	6	2	4
FY2012	6	2	4
FY2013	9	3	12
FY2014	9	3	12

Approximately 50% of the data in Table IV are accounted for in Table IIA. Approximately half of the data in Table IV are model data. The other data are observation data, which are expected to be used in data assimilation systems. Observation data are composed of radar data and operational and experimental satellite data. Vendors shall be responsible for ensuring that the radar and other observational data are delivered to the appropriate R&D HPC Subsystem.

Level II RADAR data is available on the Abilene network (Internet 2) and DC MAN.

The expected data sources for the satellite data are the current satellites still in use at the beginning of the contract possibly including METEOSAT, GOES-10-12 and NOAA-14-18 (both imagers and sounders) as well as NASA's EOS-Aqua and EOS-Terra satellites (MODIS on both and AIRS/AMSU-A on Aqua), SSM/I

(imager data) and SSM/IS (F16) on the DMSP satellites. New polar orbiting satellite data are expected to come from NOAA-N', NPP (sounder and imager in 2007), SSMIS on F17 (2005), F18 (2007), F19 (2008), F20 (2010), METOP (sounders and imager in 2005), NPOESS (sounders and imagers) C1 (2010), C2 (2011) and C3 (2013). For geostationary satellites, new sounder and imager satellite data will include GOES-N (2005), -O (2007), -P (2008) and -R (2012). All launch dates are approximate, and some need for transfer of simulated data streams of the same size as the actual data sets may occur up to a year prior to the launch dates.

Provision of the data to the appropriate Subsystems does not require the same timeliness as for NOAA's OCCS. A delay of the data transfer of up to 6 hours is acceptable. If the data is being delivered via vendor supplied communications, at the discretion of the Government, further delays of the data can be counted as downtime. It is required that no gaps in the data sets occur because of vendor supplied communications. The vendor shall notify the COTR when data are delayed or missing. If using vendor supplied communications, the vendor shall provide enough bandwidth to continue the data ingest while any data that have been lost is being re-transmitted.

C.5.4.1.1 Data Ingest requirements for Workstreams 1-2-3

While workstreams 1-3 may import data gathered through observations for model parameterizations (such as topography data), the models do not operate on observational data (i.e. there is no data "ingest" such as that associated with workstreams 4,6-9).

The INPUT directory supplied with workstreams 1 and 2 provide an idea of how much data is required to run the next segment of the model. To minimize startup costs, workstream 3 generates an idealized set of values internally. Actual data sizes required to restart the workstream 3 model are on the order of the workstream 2, CM2-HR model.

Restart and model parameter files are all located in the archive associated with workstreams 1, 2 and 3.

C.5.4.1.2 Data Ingest requirements for Workstreams 4 and 6

The data for WS 4 and 6 is received from the OCCS operational network.

Observation and model data (as described in Section C.5.2.5) are required for use by workstreams 4, 6, 7, 8 and 9. All model data will be created and will be available from the NOAA OCCS and DC MAN. Satellite data will be available from the NESDIS CEMSCS computer that is located in Suitland, Maryland (see Appendix A Figure 6) and is available on the DC MAN. The OCCS uses a small subset of satellite data. Table IV shows the projected input data quantities required by each workstream. The data ingest for workstreams 7, 8 and 9 are

similar. The data ingest for workstreams 4 and 6 are similar. About 75% of the ingest data is common among workstreams 4, 6, 7, 8 and 9.

C.5.4.1.3 Data Ingest requirements for Workstreams 7-8-9

The data for WS 7-8-9 is received from several sources. The DSRC has a satellite dish to download GOES data directly. Other data is obtained directly from their sources (including data sources and sinks within the DSRC in Boulder as well as NCAR). ETA and GFS data is received via the NWS telecommunications gateway (see Table IV). Data is also received from the AWIPS operational feed. Workstreams 7-8-9 do not have a requirement to access data saved to the data archived from the OCCS.

C.5.4.1.4 Data Ingest requirements for Workstream 5

There is no external data ingest requirement for workstream 5. All data for this workstream is generated internally.

C.5.4.2 Transfer of data to and from the R&D HSMS

The vendor shall provide all communications that interface with the R&D HSMS and all R&D HPC Subsystems and users.

The vendor shall provide any additional connectivity that exceeds GFE (see Appendix A Table VI) required by the location of a provided Subsystem. For example this requirement may be driven by the data generation on the LSC as described in the Table II in Section C.5.2.6.2 or between the post processing and analysis component and the data and archiving component.

C.5.4.3 Support for User Activities

C.5.4.3.1 User Profile

Table V shows the percentage of HPC Resources allocated for the representative workstream (WSx) according to the locality of scientific users. This profile represents the user profile expected at the start of the contract; it may change over time. "Scientific users" do not include Computer Operators, System Engineers, System Administrators and administrative users.

Table V - Users of Workstreams by Location (shown in %)

	WS1	WS2	WS3	WS4	WS5	WS6	WS7	WS8	WS9
Boulder				10%	5%	5%	65%	80%	80%
Washington				80%	80%	80%	10%	10%	
Princeton	96%	96%	96%	5%	10%				
Other	4%	4%	4%	5%	5%	15%	25%	10%	20%

User activities are grouped into the following four categories (the order listed is not indicative of relative importance):

- Software and Model Development
- Model Production
- Model Analysis
- Data Browsing

Software and model development includes activities such as code development, debugging and optimization as well as the scientific development of the models. The batch jobs submitted by this user set tend to be short, quick turnaround, but resource intensive (large PE counts, large data sets from archive or long term scratch, etc.) consistent with development of next generation models. Use of editors such as vi and XEMACS, the TotalView visual debugger and visualization tools such as GrADS, IDL, vis5D, and ncview characterize many of the interactive tools utilized by this user set. The interactive activities require good bandwidth as described below, and “LAN-like” moment to moment quality of service.

Model production is characterized by data staging and model runs generally via the submission of batch scripts. This activity includes post processing such as that for WS1-3, 5 and 8 as well as production of graphics files such as that associated with WS7-9. These production runs can be relatively few, medium to long running, self re-submitting jobs such those characterized by WS1-6 or many short jobs such as WS7-9.

Model analysis is characterized by structured and unstructured data browsing and well as interactive and/or batch jobs to produce verification statistics.

Data browsing is an activity ubiquitous to the user environment. Data browsing is more than just playing various time series of variables from an “unstructured walk” through archive datasets. The data browsing activity requires real time rendering of the data as the researcher interacts with the datasets. Manipulation of data cross sections and data range, interpolation and filtering schemes are also used in “real time” as the scientist takes an unstructured walk through a dataset as well as across datasets in multiple archive files. Significant latency in this process will degrade the experience to the point of making it useless for research.

Tools such as GrADS, IDL, vis5D, and ncview are employed to create a graphic representation of spatial and time domain data used as input to the models or created as part of model output. Those engaged in climate workstream activities such as WS1-3 and 5 tend to deal with large data files 1GB to 50GB while weather workstreams such as WS 4, 6 and 7-9 with files under 1GB. The size of the file set under analysis may be a few gigabytes to over a terabyte. Such activities require good bandwidth as described below and “LAN-like” moment-to-moment quality of service. As a further distinction, climate data browsing is

generally an “unstructured” activity (i.e. it’s not generally known in advance what data will be accessed) while weather data browsing can be viewed as “structured” (i.e. the desired datasets are known in advance).

File editing is a ubiquitous activity generally not requiring high bandwidth, but high moment-to-moment quality of service. Sufficient network bandwidth and quality of service (including high availability and low latency) shall be supplied to meet all file editing, data browsing and interactive debugging requirements concurrently. The Contractor is free to supply compression technology or other innovations compatible with the user toolset described in the following user profiles which might reduce WAN/LAN requirements and maintain the Government’s existing user experience. Testing of the data browsing, interactive debugging and file editing requirements will proceed throughout the acceptance period with periodic monitoring throughout the contract.

C.5.4.3.1.1 User Profile for WS7 – WS9

NOAA has approximately 200 scientific user accounts associated with these workstreams. Approximately 75% of those users are located at the David Skaggs Research Center (DSRC), in Boulder, CO. The remaining 25% are scattered across the country. There are approximately 40 users engaged in software and model development, 80 users engaged in data browsing roughly 30% of their time, 15 users engaged in model production and 60 users engaged in model analysis. Data browsing tends to be of the unstructured type using NCAR graphics (NCL and user applications built with NCAR graphics library), Ncview, IDL (interactively and applications), GemPAK, Vis5D, and others. Some of these data-browsing activities involve viewing single frame images while others utilize loops of fields over time as well as interactive 3D manipulations. These visualization tools are X-window based applications producing 24bit RGB images at a resolution of 1280x1024. At the time of the initial delivery of resources supporting WS7-9, a minimum of 40 simultaneous data browsing sessions producing graphics files to be animated by GFE workstations on the LAN in Boulder, CO shall be supported. The files produced will be up to 150 frames and shall be delivered to the user's GFE workstation in less than two minutes after the generation of the graphics files is completed. Additionally, a minimum of 6 simultaneous interactive TotalView debugger sessions shall be supported with LAN-like interactive use for the Boulder, CO user set.

C.5.4.3.1.2 User Profile for WS4 – WS6

NOAA has approximately 200 scientific user accounts associated with these workstreams. Approximately 80% of the users reside in the Washington DC metropolitan area. The remainder is spread across the country with some concentration at the NOAA labs in Princeton, Boulder and Miami. There are approximately 50 users engaged in software and

model development, 70 users engaged in data browsing roughly 20% of their time, 10 users engaged in model production and 70 users engaged in model analysis. Data browsing tends to be of the structured type on files 1GB or less using GrADS. GrADS is an X-window based application producing 24bit RGB images with the number of pixels equal to the resolution of the model under study. The data from WS4 (WRF-NMM) produces images at a resolution of 501x233; WS5 (GFS) produces images at a resolution of 384X190; WS6 (GSI) produces images at a resolution 768X384. GSI provides an example of the requirement at the high end. The files produced will consist of up to 150 frames and shall be delivered to the user's GFE workstation in less than two minutes after the generation of the graphics files is completed. At time of initial delivery of resources supporting WS4-6, a minimum of 30 simultaneous data browsing sessions shall be supported for the Washington, DC user set. Additionally, a minimum of 8 simultaneous interactive TotalView debugger sessions shall be supported with LAN-like interactive use for the Washington, DC user set.

NOAA will provide the connection between the DC MAN and users at the NOAA Science Center. NOAA will provide the connection between the DC MAN and users at the future NOAA facility in College Park, MD. NOAA will provide the connections from the DC MAN to both facilities during the move transition period.

C.5.4.3.1.3 User Profile for WS1 – WS3

NOAA has approximately 215 active (used in the last 90 days) scientific user accounts associated with these workstreams. Of the 96% of the Princeton users, approximately 35% are located on Princeton University's campus and are connected to NOAA's laboratory via microwave connectivity (see Section C.10.4.1). The remaining, approximately, 60% of the active accounts are users at NOAA's lab in Princeton. For these users, access to the Princeton HPCS resource is currently constrained by a 1 Gbps connection between the HPCS and the router connecting it to the users' workstations. There are approximately 30 users engaged in software and model development, 80 users engaged in data browsing roughly 25% of their time, 12 users engaged in model production and 80 users engaged in model analysis. Data browsing tends to be of the unstructured type on files 1GB to 100GB using GrADS and ncview. Ncview is an X window based application producing 24bit RGB images with the maximum number of pixels equal to the resolution of the model under study. The archive is currently mounted to each user workstation over the LAN. Thus, ncview may run as a local application on the workstation, transferring the data from the file sitting on the archive staging disk. Due to the file set size, sometimes a terabyte or more, some users must either run ncview on their workstation using the LAN mounted /archive to read the data or run ncview

from a session logged onto the Analysis cluster (see Section C.10.2.3) directly in their /archive and/or /vftmp directories, thus sending a series of RGB frames over the LAN.

There are occasions where scientists located at Princeton must transfer up to 10 TB of data from the data archive to disk storage located at the Princeton facility. These activities shall not impact the users' interactive use of the HPCS.

At time of initial delivery of resources supporting WS1-3, a minimum of 30 simultaneous data browsing sessions shall be supported for the Princeton, NJ user set. Note that within the time frame of initial delivery, the LAN supporting the Princeton, NJ user set will be provisioned to supply 10-12Gb/sec aggregate from the HPCS resources with 1Gb/sec service to user workstations. This configuration will form the baseline for the on site Princeton, NJ "user experience". Additionally, a minimum of 10 simultaneous interactive TotalView debugger sessions shall be supported with LAN-like interactive use for the Princeton, NJ user set. For information concerning the current SAN and LAN supporting these activities, see Section C.10.3.2.3.

C.5.4.4 Data Interface to the OCCS and Backup OCCS

Data must be transferred from the OCCS to the R&D HSMS. The OCCS may be in Fairmont, WV or Gaithersburg, MD at any given time (when backup and operations functions are exchanged). The backup OCCS must be able to write data to and read data from the R&D HSMS that supports workstreams 4-5-6. NOAA will provide transport for all data written from or read to these systems between these systems and the DC MAN. The vendor shall be responsible for the data transport between the HSMS supporting these functions and the DC MAN.

Each day, the NCEP OCCS produces approximately 1 TB of model output which, together with a much smaller volume of observational data, are collectively referred to as 'operational run histories'. These data are transmitted over a private telecommunications network from the operational CCS (usually Fairmont, WV) to the backup CCS (usually Gaithersburg, MD). In the event of a failover to backup operation, the direction of this dataflow is reversed. The circuit bandwidth is presently ~ 300 mbps but will be upgraded to 1 gigabit service before the end of FY05. The run histories, also referred to as the 'OCCS Write Data', are archived at Gaithersburg (in duplicate). A plan exists to begin archival of one copy at Fairmont (in lieu of two copies at Gaithersburg), but this has not yet been implemented and may not be, given the decision to transfer archival functions to the RDHPCS.

The communications circuits that carry OCCS Write Data are of high quality and reliability, with low bit error rates and very high availability. Nevertheless, data are occasionally lost and must be retransmitted when service is restored. TCP/IP

protocol ensures proper & complete transmission on a fine scale, while a layer of middleware called DBNet supervises the file-by-file queuing and transfer functions.

When circuits (or more likely the servers to which they are connected) fail to operate properly, the DBNet queues begin to build, resulting in alarms and eventually human intervention. When an outage is of short (order of minutes) duration, the system will usually catch up and recover. Longer outages result in loss of (time-sensitive) product availability which requires specific judgments and remedial actions by data flow specialists.

It also happens that data may be lost – even after successful transmission and receipt – by means of inadvertent deletion or file system failure. In these cases, a request is made for retransmission (and sometimes recreation) of the missing data.

C.5.5 IT Security Requirements

To assure an adequate level of protection for in-house or commercially maintained IT systems, NOAA maintains all systems consistent with government-wide laws and regulations. The Office of Management and Budget (OMB) Circular A-130 requires all federal agencies to plan for the security of all IT systems throughout their life cycle. OMB Circular A-130 requires agencies to implement and maintain a program to assure that adequate security is provided for all agency information collected, processed, transmitted, stored, or disseminated in general support systems and major application. It further directs each agency to implement policies, standards, and procedures that are consistent with government-wide policies, standards, and procedures issued by the Office of Management and Budget (OMB) and Department of Commerce (DOC). OMB Circular A-130 also establishes a minimum set of controls to be included in Federal automated information security programs and assigns Federal agency responsibilities for the security of automated information. At a minimum, agency programs will include the following controls in their general support systems and major applications:

- Assignment of Responsibility for Security.
- System Security Plan consistent with guidance issued by the National Institute of Standards and Technology (NIST) to include:
 - Rules of Behavior/Application Rules.
 - Training.
 - Personnel Controls.
 - Incident Response Capability.
 - Continuity of Support/Contingency Planning.
 - Technical Security/Controls.
 - System Interconnection/Information Sharing.
 - Public Access Controls as required.
- Scheduled Review of Security Controls commensurate with the acceptable level of risk to the system.
- Authorize Processing (Certification and Accreditation).

The Federal Information Security Management Act of 2002 (FISMA), (HR 2548 E-Government Act, TITLE III – INFORMATION SECURITY, SEC. 301. INFORMATION SECURITY), addresses the program management and evaluation aspects of IT security.

Additionally, NOAA follows the US DOC Manual of Security Policies and Procedures, the US DOC IT Security Program Policy and Minimum Implementation Standards, as well as NOAA OCIO Policies and Directives. NOAA 212-1300 IT Security Policy provides specific guidance on Patch Management and Risk Management controls to be used for all NOAA systems.

More information can be found at:

- <http://www.cio.noaa.gov> (NOAA CIO web-site)
- <https://www.csp.noaa.gov> (NOAA IT Security Office web-site)
- <http://www.osc.doc.gov/cio/oipr/ITSEC/ITSECDOC1.HTM> (DOC IT Security Program Office web-site)
- <http://csrc.nist.gov/policies/> (NIST Federal Requirements)
- <http://csrc.nist.gov/publications/nistpubs/index.html> (NIST Guidelines)

The Government reserves the right to augment or update the list of compliance documents. The Contractor shall be responsible to assure that all of the Government Security requirements are met for this system. [The Contractor shall provide Configuration reports, monitoring reports, security assessments, and compliance matrices upon request to the COTR or Government ITSO. The proposal shall clearly indicate the methodology and personnel that will perform all security tasks.](#)

The designated System Owner (SO) will be:
William T. Turnbull, NOAA Deputy CIO & Director, NOAA HPCC Office

The IT [System](#) Security Officer (ITSSO) will be [designated by the Government.](#)

These Government officials may delegate aspects of system resource documentation and management to the Contractor, but will retain full authority over system resources and information security. The Government authorizes the use of system computer resources for Contractor performance of the effort required by the statement of work of this contract. Under no circumstances is a Contractor permitted to make any use of organization computer equipment or supplies for purposes other than performance on this contract.

The offeror shall provide a plan that describes its IT security program. The plan shall address the security measures and program safeguards, which will be provided to ensure that all information systems and resources acquired and utilized in the performance of the contract by Contractor and Sub-Contractor personnel:

- Operate effectively and accurately

- Are protected from unauthorized alteration, disclosure, or misuse of information processed, stored, or transmitted
- Can maintain the continuity of IT support for organization missions, programs, and function
- Incorporate management, operational, and technical controls sufficient to provide cost-effective assurance of the system's integrity and accuracy
- Have appropriate technical, personnel, administrative, environmental, and access safeguards
- Notify Government of any and all vulnerabilities found.

This plan will be included in any resulting contract for Contractor compliance.

In order to maintain a consistent security posture across potentially multiple sites, the Contractor shall be required to implement DOC/NOAA equivalent Patch Management, Change Management, and Perimeter Protection procedures for each processing site. These procedures shall be documented and maintained as part of the required NIST SP 800-37 Certification and Accreditation (C&A) package. All IT equipment delivered as part of this procurement is to be considered a Government computing resource, regardless of its location or actual owner. The Government will initiate the C&A package for this system and then transition maintenance of the security documentation to the Contractor after delivery and acceptance. The Contractor may also elect to employ equipment which is not part of the delivery for use in remotely diagnosing, monitoring, or managing the delivered system. If this Contractor-owned equipment has any special access or trust relationship with the delivered system, then the Contractor shall initiate and maintain any appropriate C&A documentation for this equipment as a separate Government system, as specified in Section H.16. Within two weeks of contract award, the Contractor shall be responsible for ensuring that all aspects of the C&A process are completed in a timely manner.

The R&D HPCS will be managed as a single system regardless of physical location. Multiple-site solutions might implement shared resources (e.g., public Internet infrastructure) for inter-site communication. Special care must be taken to secure inter-site communication over any untrusted network link, and mitigate the risks associated with the inherent trust relationships required for full R&D HPCS integration. [All remote access to the HPCS from untrusted networks shall require use of a secure multifactor authentication mechanism which minimizes the risk of stolen-credential attack. The proposed solution should initially make use of the token based authentication mechanisms at Boulder and Princeton until a common authentication framework becomes available through Homeland Security Presidential Directive 12.](#) While the Government may provide some measure of physical or network perimeter protection for the system, the Contractor shall be responsible for implementation of IT Security for the delivered solution in accordance with DOC and NOAA policies. In instances where communication or interaction is required across system boundaries, standard (NIST SP 800-47) interconnection agreements will define the security responsibilities and procedures for all parties involved. [System perimeter protection controls such as firewalls may be bypassed in](#)

cases where performance needs can be met safely using other security controls. Doing so requires close coordination with the System Owner, the ITSSO and COTR of the connecting Systems; and must be documented as part of the System Security Plan per NIST SP800-43, "Security Guide for Interconnecting Information Technology Systems" (<http://csrc.nist.gov/publications/nistpubs/800-47/sp800-47.pdf>). Trust relationships effectively extend the scope of vulnerability beyond one System's realm of control, and as such should be chosen judiciously and monitored closely. Special access to the R&D HPCS, beyond that given to remote users in DOC policy, will require additional security controls in order to mitigate the increased risk. Vendors are strongly encouraged to review the June 30, 2005 update to DOC's "IT Security Program Policy and Minimum Implementations Standards" (http://www.osec.doc.gov/cio/ITSIT/DOC_IT_SEC_pgm.htm) document for additional guidance and required security controls. All solutions shall comply with all DOC and NOAA security policies. Industry best practices will be used to increase the security posture of the system beyond the DOC and NOAA policy requirements. NIST guidelines/practices are the preferred industry [guidelines](#)/best practices.

The system shall:

- Include a mechanism to require users to uniquely identify themselves to the system before beginning to perform any other actions that the system is expected to mediate.
- Be able to maintain authentication data that includes information for verifying the identity of individual users (e.g., passwords)
- Protect authentication data so that it cannot be accessed by any unauthorized user.
- Be able to enforce individual accountability by providing the capability to uniquely identify each individual computer system user.
- Raise alarms when attempts are made to guess the authentication data either inadvertently or deliberately (based on a number of incorrect password attempts).

For further information see U.S. Department of Commerce IT Security Program Policy Section: 16

The system shall use identification and authorization data to determine user access to information. The system shall be able to define and control access between subjects and objects in the computer system. The enforcement mechanism (e.g., self/group public controls, access control lists, roles) shall allow users to specify and control sharing of those objects by other users, or defined groups of users, or by both, and shall provide controls to limit propagation of access rights. The discretionary access control mechanism shall, either by explicit user action or by default, provide that objects are protected from unauthorized access. These access controls shall be capable of including or excluding access to the granularity of a single user. Access permission to an object by users not already possessing access permission shall be assigned by only authorized users.

For further information see U.S. Department of Commerce IT Security Program Policy Section: 17

The system shall be able to create, maintain, and protect from modification or unauthorized access or destruction of an audit trail of accesses to the objects it protects. The audit data shall be protected so that read access to it is limited to those who are authorized.

The system shall be able to record the following types of events: use of identification and authentication mechanisms, introduction of objects into a user's address space (e.g., file open, program initiation), deletion of objects, and actions taken by computer operators and system administrators and other security relevant events. The system shall also be able to audit any override of human-readable output markings.

For each recorded event, the audit record shall be able to identify the date and time of the event, user, type of event, and success or failure of the event. For identification and authentication events, the origin of request (e.g., terminal ID) shall be included in the audit record. The system administrator shall be able to selectively audit the actions of any one or more users based on individual identity and/or object label.

For further information see U.S. Department of Commerce IT Security Program Policy Section: 18

For planning purposes, the R&D HPCS system will be designated as NOAA Mission Critical, with a FIPS 199 security categorization of {(confidentiality, low), (integrity, moderate), (availability, moderate)}. The resulting moderate-impact level designation may then be used to select recommended controls from NIST SP800-53 ("Recommended Security Controls for Federal Information Systems"). However, the final system categorization will be determined through development and approval of the System Security Plan (SSP).

The following clauses are reprinted from FAR 52.239-1, Privacy or Security Safeguards.

- (a) The Contractor shall not publish or disclose in any manner, without the contracting officer's written consent, the details of any safeguards either designed or developed by the Contractor under this contract or otherwise provided by the Government.
- (b) To the extent required to carry out a program of inspection to safeguard against threats and hazards to the security, integrity, and confidentiality of government data, the Contractor shall afford the Government access to the Contractor's facilities, installation, technical capabilities, operations, documentation, records, and databases.

- (c) If new or unanticipated threats or hazards are discovered by either the Government or the Contractor, or if existing safeguards have ceased to function, the discoverer shall immediately bring the situation to the attention of the other party. The Contractor shall provide an analysis of the new threat, hazard, or vulnerability and recommend possible fixes or safeguards.

The following Commerce Acquisition Regulation (CAR) security clause is applicable to this contract:

Security Processing for Contractor/SubContractor Personnel Working on a Department of Commerce Site (Low and Moderate Risk Contracts) (CAR 1352.237-72 (MAR 2000))

Security Processing Requirements

(1) U.S. Citizens Working on DoC Site

All Contractor (and subContractor) personnel proposed to work on the premises of a Department of Commerce site for 180 days or more shall undergo security processing by the Department's Office of Security (OSY) to be eligible to work on the site.

(2) Foreign Nationals (Non-U.S. Citizens)

Regardless of anticipated length of on-site work, all foreign nationals to be employed under this contract shall:

a. Have legal visa status with the Immigration and Naturalization Service (INS); Have advance approval from the serving Security Officer in consultation with the Office of Security.

b. Submittal Requirements – U.S. Citizens

(1) Duration of Onsite Work: 180 to 364 days (between 6 months and 1 year)

For individuals who will be performing work on a DoC site between 180 and 364 days, the Department will perform a Special Agreement Check (SAC). The scope of the SAC will include checks of the Security/Suitability Investigations Index (SII), other agency files (INVA), Defense Clearance Investigations Index (DCII), FBI Fingerprint (FBIF), and the FBI Information Management Division (FBIN).

The Contractor shall complete and submit the following form to the Contracting Officer's Technical Representative (COTR):

Form FD-258 (Fingerprint Chart)

Copies of this Form can be obtained from the COTR. Upon receipt of the FD-25, the COTR will complete form OFI 86C (Special Agreement Check) and will forward both to the operating unit Security Officer. The Security Officer will advise the COTR whether work can commence prior to suitability determination, based on the specifics of the situation. The COTR will notify the Contractor of an approved contract start date as well as favorable findings of the suitability determination.

(2) Duration of Onsite Work: 365 days (1 year) or more
Individuals proposed to perform work on a DoC site for 1 year (365 days) or more are required to have a NACI check (National Agency Check Plus Written Inquires).

The Contractor shall complete and submit the following forms to the Contracting Officer's Technical Representative (COTR):

Standard Form 85P (SF-85P, Questionnaire for Public Trust Positions), and FD-258 (Fingerprint Chart).

Copies of these Forms can be obtained from the COTR. Upon receipt of the required forms, the COTR will forward the forms to the operating unit Security Officer. The Security Officer will advise the COTR whether work can commence prior to suitability determination, based on the specifics of the situation. The COTR will notify the Contractor of an approved contract start date as well as favorable findings of the suitability determination

c. Submittal Requirements – Foreign Nationals

All Foreign nationals proposed to work on a DoC site will be subject to a Special Agreement Check (SAC) to determine whether the foreign national has official legal status in the United States.

The Contractor shall submit the following forms to the COTR for all foreign nationals proposed to work on a DoC site:

FD-258 (Fingerprint Chart)

Form OFI 86C (Special Agreement Check) with signature authorization for release of information
Copies of these Forms can be obtained from the COTR. Upon receipt of the required forms, the COTR will forward the forms to the operating unit Security Officer. The COTR will notify the Contractor of favorable findings and will notify the Contractor regarding an approved date to commence work under the contract.

d. Suitability Updates

Any individual (including foreign nationals) processed on the form OFI-86C (Special Agreement Check) who stays on the contract over 364 days will be required to have a NACI complete suitability check to stay on the job site.

e. Notification of Disqualifying Information

If OSY receives disqualifying information on a contract employee, the Contractor, upon notification of such by the Contracting Officer, must immediately remove the employee from duties which require access to DoC facilities.

Individuals may be barred from working on the premises of a facility for any of the following:

- (1) Conviction of a felony or a crime of violence or of a misdemeanor involving moral turpitude.
- (2) Falsification of information entered in security screening forms or on other documents submitted to the Department.
- (3) Improper conduct once performing on the contract, including criminal, infamous, dishonest, immoral, or notoriously disgraceful conduct or other conduct prejudicial to the Government regardless of whether the conduct directly related to the contract.
- (4) Any behavior judged to pose a potential threat to departmental personnel or property.

Failure to comply with the requirements may result in termination of the contract. Compliance with these requirements shall not be construed as providing a contract employee clearance to have access to classified information.

C.5.6 Facilities Requirements

The Contractor shall provide a facilities proposal for housing the proposed equipment to meet the Government's requirements over the life of the contract. Contractors may use any of the following to meet the Government's needs: Any of the available Government facilities (as described in Section C.11), either as is or with modifications; Contractor-provided facilities; or a combination of Government facilities and Contractor facilities. The Contractor shall provide details of their proposed facilities solution in the Facilities Proposal that they include in their submission.

C.5.6.1 Government-Provided Facilities

Contractors may use the Government-furnished facilities offered at the NOAA locations in Boulder, CO, Princeton, NJ, and Greenbelt, MD, which are described in Section C.11. Section C.11.11 indicates the Government-supported resources that are projected to be available, beginning at the times indicated and provided under the assumptions indicated in that section. If Contractors require additional facility resources at these locations beyond what is indicated, they shall provide a modification plan in their proposal that indicates proposed changes, with schedules, for the work to be done. The modification plan shall be in accordance with site restrictions.

Contractor use of the Government-supported resources assumes the following conditions apply:

- The Government is responsible for delivering the indicated facilities resources (although the Contractor will be responsible for paying power utility costs in accordance with Section H.19). Because of this, any lost system time caused by environmental outages or facility failures will be recorded as null time for availability calculations.
- The Contractor shall provide the funding for any facility modifications (e.g., new electrical equipment, plumbing, additional air conditioning equipment) needed to install and support the proposed systems, and to obtain facility resources in excess of what the Government makes available as indicated in Section C.11.11. The Contractor shall be responsible for carrying out the facility modifications at facilities such as the Princeton facility where the Contractor is permitted to perform the build-out; for these facilities, the Contractor shall include a description of these modifications, including schedules, as part of its Facility Proposal.
- Any new facility equipment (e.g., air handlers, chillers, UPSs, or switchgear) or building modifications will become the property of the owner of the facility, with the exception of UPS equipment at PRTN. All facility equipment (e.g. PDUs, air handlers) will be maintained by the owner of the facility, unless otherwise negotiated with the Contractor. However, UPS equipment at PRTN, either GFE or vendor-owned, that is used to support systems under this contract shall be maintained by the Contractor. All other UPS equipment will be maintained by the Government or the facility owner.
- If the Contractor chooses to use the computer-conditioned space at the Greenbelt site, hereafter referred to as GRBLT site, it shall include a per annum cost of \$130,000 in its price proposal for use of the space, beginning in October 2006. This cost shall be in addition to the electric utility costs as prescribed in Section C.11.3.2.
- For the GSA-managed facilities, BLDR-1 and BLDR-2, in which the any facility modifications must be paid to the facility manager by the Government prior to the work being contracted for and performed by the facility manager (e.g., GSA), the Government shall be fully reimbursed by the Contractor for the work to be performed through credits against the contract that do not extend beyond the fiscal year in which the expenditures are made by the Government.

For any facility modifications to be made, the Contractor shall follow the following procedures that are specific to each of the indicated sites:

Boulder, CO

The David Skaggs Research Center (DSRC) is owned by the Government and operated by the General Services Administration (GSA). Because of this, all building modifications are subject to GSA Regulations, thereby requiring that GSA contractors carry out these modifications. Therefore, any facilities modifications will

require that funds be held back from the contract for transfer to GSA for any changes to be made to GSA properties.

Close coordination with the DSRC landlord, U.S. General Services Administration (GSA), will be required if any building modifications are required. Costs will include normal GSA overhead charges plus project management oversight. All modifications must meet standard building codes as dictated by GSA. The Contractor is advised that, because of the long implementation delays for large projects (see table below), facility modifications, at least during the first year of the contract, shall be limited to small projects such as power receptacle changes. Large facility modifications will require substantial lead-time to complete.

Examples of Boulder Facility Modifications (Estimations ONLY)

Small Facility Modifications		
Modification	Estimated Cost¹	Estimated Time Required
Power Receptacle Relocation (per receptacle)	\$1500	60 Days
Power Receptacle Alteration (per receptacle)	\$2000	60-90 Days
Large Facility Modifications		
Additional Chiller and Cooling Tower	\$1.2 Million	8 to 10 Months
Additional Emergency Generator	\$1 Million	8 to 10 Months
¹ Note: Costs include installation, support equipment and associated building modifications. Costs do not include environmental studies, feasibility studies, building impact surveys or GSA overhead and program management fees.		

All facility infrastructure modifications and equipment at DSRC will become the property of the Government and will be fully maintained by the Government in accordance with manufacturers' maintenance schedules. Normal Government contracting procedures can be expected for any modifications.

The Contractor shall be responsible for any costs that exceed the available GFE power resources, utility power costs and associated maintenance costs (both routine and emergency) for electric usage. These costs shall be paid for via monthly contract deductions.

Maintenance costs for equipment (both routine and emergency) that is installed to enable power resources to exceed the GFE offered amount identified in Section C.11.11, shall be the responsibility of the Contractor. These costs shall be paid via monthly contract deductions.

Princeton, NJ

Because the Princeton buildings are owned by Princeton University, all construction shall be performed and paid for by the Contractor and is subject to all applicable building codes and approvals for the work from Plainsboro Township and Princeton University. In addition, if these facility modifications cause any changes to the exterior of the building, these modifications shall be in accordance with Princeton University's Princeton Forrestal Center Design and Development Criteria, which includes approval by the Design Review Committee.

If the Contractor's projected facility requirements exceed the available GFE power and/or cooling, the Contractor shall be responsible for implementing any necessary enhancements to the facility infrastructure to support this requirement. In projecting the needed enhancements, the Contractor may assume the following:

- 1) Up to 3,000 kW of service to the GFDL facility, split between two 1,500 kVA transformers, one in the Main Building and the other in the Computer Building.
- 2) By spring 2006, PRTN will have completed the installation of a chilled water loop capable of providing sufficient chilled water flow to permit the running of any combination of its three chillers. The designs and specifications for this are available from the Contracting Officer upon request.

Also, in order for the Government to upgrade the campus substation to meet the required power indicated in assumption (1.) above, the winning Contractor shall document and provide to the Government its anticipated power loading requirement for its PRTN system within one month after contract award. This information will permit the Government to upgrade the electric service to PRTN through discussions with its landlord and local electric utility provider.

Given that the current HPC contract with Raytheon ends September 30, 2006, the Contractor shall assume, for the purpose of this evaluation, that the current HPC equipment will be powered down at that time and will not consume any power or cooling resources after that date. Further, the Contractor shall assume, for the purpose of this evaluation, that all non-GFE equipment in the current contract that is not to be used in the new R&D contract will be removed from the Computer Room within 30 days after this date. Any deviations from this plan shall be clearly identified in the transition plan and be agreed to by the Government. The successful Contractor is expected to work with Raytheon to coordinate this transition; however the Contractor shall keep the Government COTR apprised of all dialogs and issues.

Greenbelt, MD

NASA/Goddard Space Flight Center (GSFC), located in Greenbelt, Maryland, is one of several NASA centers in the United States. NOAA and GSFC recently signed a Memorandum of Understanding under which GSFC has agreed to provide NOAA with computer-conditioned space. This agreement takes effect on April 6, 2005, and runs through April 6, 2010; it can be extended beyond this end date by mutual

agreement of NOAA and GSFC. NOAA is offering this computer-conditioned space at the GRBLT site for Contractors to consider as they prepare their proposals in response to this RFP.

NOAA and GSFC will carry out the site preparation into two phases. Phase I will involve the construction of the computer room from the available office space on the ground floor of GSFC Building 28 to produce approximately 5500 square feet of contiguous raised floor space. NASA will make available up to 260 tons of cooling and 900 KVA of power to the computer room.

For Phase II of the site preparation, NOAA and GSFC's Facilities Management Division (FMD) will work with the winning Contractor in developing a detailed site plan, using the Contractor's facility proposal and transition plan. Post-award changes to the computer room, subsequent to the Phase I effort above, may include, depending on the selected Contractor's specific equipment, purchase and installation of CRAC units and PDUs, purchase and installation of new floor tiles, and modifications to accommodate above-standard equipment height, etc. GSFC's FMD will review and approve all proposed changes to the offered space. This approval process is estimated to require several weeks to complete, from the date that the proposal is submitted by the Contractor, through the NOAA Contracting Officer.

NOAA will provide the funds outside of the R&D contract to cover the cost of the Phase I site preparation. The cost of the Phase II site preparation shall be included in the Contractor's cost proposal. All costs associated with the Phase II site preparation at GFSC shall be the Contractor's responsibility. Accordingly, the Contractor shall adjust its payment profile in FY2007 to correspond to its price proposal for the Phase II site preparation. The reduction in FY2007 payments shall be amortized in equal amounts over a six-month period commencing in October 2006.

C.5.6.2 Contractor-Provided Facilities

The Contractor may propose to use facilities other than the Government-provided facilities that are described in Section 11, Appendix B. Any Contractor-provided facilities shall meet the following conditions:

- Contractors are responsible for providing all facilities resources, including floor space, utilities, facilities maintenance, janitorial services, etc. Therefore, any lost system time caused by environmental outages (such as loss of power, cooling, etc.) or facility failures will be recorded as downtime for availability calculations.
- Contractors may use any Unrestricted GFE provided in Appendix C, Section C.12, but are responsible for all shipping costs, and for the shipment both to the Contractor site and back as directed by the Government.
- Designated Government personnel must have access to the Contractor facility on an as-needed basis, subject to reasonable controls.

C.5.6.3 Office space, workstations, LAN connectivity, phone service and office equipment for on-site contract personnel

The following policies will be in force for any Contract personnel to be located at Government sites.

The Government will provide LAN connectivity to all Contractor personnel located at its sites. The Contractor shall provide workstations, phone service, and office furniture for all of its personnel located at Government sites, although the Government may, for its convenience, choose to provide this on a case-by-case basis. To assure that IT security policies are maintained, the Government will require Contractors to install, operate, and maintain their workstations on the LAN in a manner that is consistent with Government IT security policies. [The Government will provide access to fax and reproduction machines for reasonable use.](#)

Boulder, CO

The Government will supply two (2) 150-sq. ft. offices and approximately 100 square feet of lab space for use by the Contractor. [The Government will provide phone service.](#)

Washington, DC

NCEP can accommodate two on-site personnel at the World Weather Building and its annexes since office space is extremely limited.

Greenbelt, MD

[No office or lab space is provided. Contractors requiring office and/or lab space must use space within the provided room. Contractors will have access to a party-line phone within the computer room for making and receiving calls.](#)

Princeton, NJ

Effective October 2006, the Government will provide two offices for use by Contractor's personnel who are responsible for system operation. These two offices will be occupied by the current Raytheon system support staff until the end of the contract on September 30, 2006. They are the large room adjacent to the Operators Room and an office across the hall. If the R&D Contractor requires additional office space and/or storage space beginning in October 2006 or requires office and/or storage space prior to October 1, 2006, it may construct offices on the hard pan in the southwest corner of the Computer Room as part of its facility proposal. It may construct these rooms in of the hard pan space that was previously planned for offices as indicated in Figure 13. If the Contractor does not choose to use this hard pan for this purpose, it may use this non-raised floor for installation of equipment that does not require raised floor.

In addition to the two offices provided for personnel responsible for system operation, the Government will attempt to provide additional office space to other Contractor personnel on an as-needed and as-available basis.

The Government will provide a party line phone within the computer room that the Contractor may use for receiving incoming calls and making outgoing calls; however, Contractor personnel will not have voicemail capabilities on this party line. The Contractor shall provide telephone service to the provided GFE office space.

C.6 Reliability and Availability Requirements

C.6.1 Reliability and Availability

The HPCS offered shall continue NOAA's historically high utilization of its computing resources. Reliability, availability, and Contractor support are considered fundamental aspects of the HPCS.

C.6.1.1 Reliability

Downtime (as defined in Section C.13.2, Appendix D) will be used in the determination of the actual System Life Throughput. Periods of Remedial and Preventive Maintenance count as downtime. Null time will not be counted as downtime. Null time is that period of time when the workload cannot be accomplished due to environmental failure at a Government-furnished site, such as loss of electric power or cooling, or recovery from environmental failure. Downtime for each HPCS component is based on the fraction of the resources available for that component's workload. It is arrived at through consultation between the Government and the Contractor, and ultimately determined by the Government. Downtime is accumulated on the HPCS if the Government is not able to execute its workload when access to the scientific data via the storage or archiving components is unavailable.

A component's downtime shall commence from the time the Government or Contractor determines that a downtime condition has occurred. Upon the discovery of a problem with the system, the Government will contact the Contractor at the designated point of contact (see Section C.7.1). The Contractor should not solely rely on the notification by the Government to determine a downtime condition. The Government, in order to monitor the system, will have at least one operator on duty, 7x24x365. The Government will log the problem which will be updated with information that is provided by the Contractor. The information to be entered into the log will be determined by the Government.

A component's downtime shall exclude any time in which the Government denies the Contractor maintenance personnel access to the malfunctioning hardware and software. (Routine procedures for entry to a facility shall not be construed as denial of access.)

A component's downtime shall end when it is returned to the Government in operable condition as determined by the Government, ready to perform all of the workload.

Preventive Maintenance (PM) is to be completed at times determined by the Government.

The testing and installation of every major operating system release installed at the request of the Government and one (1) minor operating system release installed at the request of the Government during any annual period will count as downtime. The installation of security patches shall count as downtime. Preparation for and execution of post-upgrade LTDs, including any benchmark development by the Contractor, associated with the agreed-upon upgrades proposed at contract award, will count as downtime.

The Government, at its discretion, may provide a series of carefully monitored jobstreams that may include up to twenty-five (25) individual batch jobs in order to verify the reliability of the system. These jobstreams will be monitored for end-to-end success. Complete or partial failure of any intermediate step will result in declaring the jobstream to have failed. Success ratio for a jobstream will be determined as the number of successful jobstreams divided by the number of attempted jobstreams in a 30-day period per subsystem. If multiple types of jobstreams are being monitored on a subsystem, then the aggregate of their success ratios will be computed by weighting by the total amount of CPU-hours each jobstream should use over a 30-day period. The Government requires that ratio be greater than or equal to the proposed availability. This will be evaluated monthly. Failures due to null time, application errors, or data errors will not be included in the numerator or denominator of the above ratio. The Government is responsible for ensuring that the applications function properly on the subsystem. The Government may, at its discretion, choose to accept tools, support, and advice that will make the jobstreams more reliable.

C.6.1.2 Availability

All components of the HPCS must perform as an integrated system to provide the Government with at least 96% availability. These components include not only those typically associated with a high performance computing system but also any LAN/WAN components provided as part of the HPCS. Additionally, the Government requires 99% availability for access to its scientific data. Scientific data archives may be on the Long-Term Scratch component or the HSMS component, depending on the overall architecture associated with a given workstream. The Government's data are critical assets associated with the HPCS. With the large amount of data associated with the HPCS, the Government requires the highest levels of data integrity.

Although availability will be monitored daily, it will be measured on a monthly basis. The Contractor shall establish and maintain accurate system monitoring with

associated output data to support system availability measurements. Monitoring tools shall be made available to the Government. Throughput shall be determined by relating system availability to the workstreams targeted to run on that system. For example, if Subsystem A typically runs workstreams 1 – 2, workstreams 3 – 9 will not be used to determine the amount of “lost” throughput, if any.

Availability shall be determined by computing the ratio of total computation socket hours available for execution of R&D jobs to the total computation socket hours each month, excluding Null Time. The time a computation socket is available for execution is determined by subtracting socket downtime from wall clock time. Spare sockets can be included in the computation pool to reduce downtime. However, downtime accumulates until the spare sockets are made available for job execution. Accumulated computational cycles (in CPU-hours) that are lost when jobs are lost due to component failure or component reboot will not count toward the system-life throughput calculation. If the accounting software cannot report the accumulated computational cycles for each active job at the time of failure, it will be assumed that four (4) [socket](#)-hours were lost for all socket(s) on the failed system.

System Life Throughput (SLT), for a given workstream i , can be calculated by the following equation:

$$SLT_i = \sum_{j,k} \frac{T_{i,j,k} A_{i,j,k}}{B_{i,j,k}}$$

SLT	= System Life Throughput
T	= Total Wall-clock Time during system configuration k
A	= Availability
B	= Workstream Instance (Benchmark) Throughput Time
i	= WS number
j	= WS instance
k	= system configuration period

In words, the SLT for workstream i is the Total Wall-clock Time for a given configuration k multiplied by the availability for that configuration and divided by the Workstream Instance Throughput Time summed over all workstream instances j and system configurations k . The summation over instances is done because the target IT architecture for the workstream may contain heterogeneous socket types. As an example, some instances of a workstream may run with a given socket revision (and associated availability) while other instances run on a different socket revision (and its associated availability).

An example of a term in the System Life Throughput sum for a sample workstream can be demonstrated for FY2001-FY2003. This calculation is based on actual calendar days beginning October 1, 2000. For example, if the proposed LSC throughput benchmark execution time is 10,800 seconds initially, and is upgraded to 7,200 seconds on October 1, 2001, and a 98% availability is proposed through the three-year period, the system life throughput SLT component for the instance j is given below (Note: there are 86,400 seconds in a day).

$$SLT^{sample}_{i,j} = \frac{((365 \frac{days}{year})(86400 \frac{s}{day}))(0.98)}{10,800 \frac{s}{benchmark}} + \frac{(2(365 \frac{days}{year})(86400 \frac{s}{day}))(0.98)}{7200 \frac{s}{benchmark}}$$

$$SLT^{sample}_{i,j} = 11446.4 benchmarks$$

Note that the throughput value at upgrade (7200) may be from running the instance on faster hardware or running the instance on more sockets of the original configuration. This is because with the exception of WS5, the number of instances of a workstream is constant over time. Thus if at upgrade additional sockets are added to an existing configuration, then the fixed number of workstream instances targeted for the upgraded configuration must be redistributed to cover old and new components of the upgraded configuration.

For example, consider an initial system configuration A1 running the 8 instances of WS1 and 4 instances of WS3. At upgrade the system configuration for WS1 and 3 becomes A1 + B1. The 8 instances of WS1 and 4 instances of WS3 must now be distributed across A1 + B1. The Contractor might run 3 instances of WS1 and 1 instance of WS3 on A1 at higher core counts than the initial installation and 5 instances of WS1 and 3 instances of WS3 on the (presumably) more capable B1. Or as another example, the Contractor may choose to run all 8 instances of WS1 on A1 or B1 and the 4 instances of WS3 on the other machine. Specifically, it is incorrect to run 8 instances of WS1 and 4 instances of WS3 on A1 and 8 instances of WS1 and 4 instances of WS3 on B1 to determine the SLT component at upgrade. At all points in time, there are only 8 WS1 instances and 4 WS3 instances running across the entire configuration.

In calculating the offered SLT, the Contractor shall count all processing cores proposed for model computation and analysis provided those cores are configured in such a way as they may be part of a socket set running an instance of a workstream. Because the post-processing of WS 1-3 will be running concurrently with models themselves, processor cores proposed for WS 1-3 post-processing shall be excluded from the SLT calculation. Additionally, processor cores dedicated to file systems or other I/O activities shall be excluded from the SLT calculation.

For the RFP response, the Contractor shall provide a spreadsheet showing the explicit SLT calculation. The Contractor shall clearly identify the processor core count and target system configuration for each SLT component. Further, the Contractor shall identify the purpose and configuration of any processor cores not used for the SLT calculation, but present within the proposed configuration. This identification includes, but is not limited to, processor cores which for whatever reason are left idle for the purposes of the throughput benchmark. In addition, the Contractor shall provide the rationale for the processor core configuration used for the throughput benchmark (e.g. 6 of the 100 processor cores on each node offered

for WS1 are left idle during the throughput benchmark because of benchmark model decomposition and configuration tiling limitations).

Proposed throughput benchmark performance levels will be combined with the proposed availability level (98% is shown in the example above) to determine a measure of overall proposed System Life Throughput for the workstream. The actual throughput will be measured on a daily basis by monitoring the availability of the components associated with the given workstream. All performance levels shall be met for each measurement of actual throughput regardless of past delivery of suites.

Although system throughput shall be monitored daily, the accumulated System Life Throughput of the workstream shall be calculated monthly. At the sole discretion of the Government, shortfalls in throughput, when compared to the proposed throughput, shall be made up with either the delivery of additional equipment installed at no cost to the Government, or accrual of downtime credit. Should the Government elect the delivery of additional equipment to satisfy throughput shortfalls, the Government will calculate the duration of the compensation using the demonstrated benchmark performance on the upgraded HPC Subsystem. It is the Government's goal to meet or exceed the total System Life Throughput by the end of the contract in a manner that does not require frequent disruptions, front-loading of cycles, or back-loading of cycles. *As such, the Government will schedule compensatory upgrades with planned upgrades, such as the mid-life upgrade, or every 6 months. The Government may make special arrangements when there is less than 6 months left in the contract.* Should the Government elect downtime credits in lieu of additional equipment, downtime credits will be deducted from the most recent invoice. Downtime credits shall accrue on all components of the system(s).

At the sole discretion of the Government, shortfalls in proposed availability or performance of non-computational components shall be rectified with the delivery of additional equipment or engineering, at no cost to the Government, or accrual of downtime credit. To better reflect NOAA's computational needs over time, changes in the HPCS workstream benchmarks shall be made by mutual agreement between the Government and the Contractor throughout the life of the HPCS. If the parties are unable to reach mutual agreement, the Government reserves unilateral right to determine changes to the HPCS workstream benchmarks and the Contractor agrees that it will comply subject to the Disputes Clause.

The Government requires a credible method of maintaining system availability. The Government expects that fail-over and high reliability components will be used. The Government requires sufficient power during environmental failure to gracefully shut down all components of the HPCS. Further, the Government requires adequate power conditioning to insulate the System from power spikes and sags. It is expected that uninterruptible power supplies (UPSs) and power

distribution units (PDUs) will be needed for all components of the HPCS to meet these requirements.

C.7 Support Services Requirements

C.7.1 Support

The Contractor shall provide the Government with a designated point of contact to request maintenance. The Contractor shall maintain escalation procedures that allow the Government round-the-clock telephone contact with knowledgeable Contractor staff should the designated point of contact be unavailable.

For each provided HPC Subsystem, the Government requires comprehensive support in order to meet the 96% availability requirement. Support professionals can consist of systems analysts, hardware engineers, and applications analysts. Contractor system analysts shall work closely with Government system administrators. Government system administrators will have root access to the HPCS computing platforms. The Government requires a senior applications analyst to be located at each of the primary user sites for WS1 – WS9. Senior Application Analysts shall be available, in person, during business hours. The primary user sites are described in the User Profile Section C.5.4.3.1. NOAA expects to provide offices for the senior analysts at its facilities (see Section C.5.6.3 for more information). It is anticipated that additional senior application analysts may be required during transitions to new technologies (see Section C.9.5).

The Government requires the Contractor to maintain a current, itemized list of all Contractor-supplied hardware and software items in printable electronic form. The listing shall be updated whenever a system upgrade or engineering change proposal results in a change to the system configuration or delivery of additional equipment. The listing is to be submitted to the COTR upon request.

C.7.2 Training

Training shall be provided for NOAA computer specialists and operators in:

- system administration and tuning
- hardware operation and system overview

Training shall be provided for a no more than 100 NOAA applications programmers in:

- application and shell programming
- programming languages and tools
- HSMS software user interface
- Optimization

Training shall commence within 60 days from the date of system acceptance. If any significant (as determined by the Government) software or architectural changes are made to the system(s), a new training session(s) shall be scheduled within 60 days.

The Contractor shall provide the Government with a list of additional potential training topics. Training shall be collocated with users to the greatest extent possible. On-line training with access to experts will be considered.

The Government desires to begin training when early access to systems similar to those proposed for the HPCS is granted. See Section C.5.1.2 for more information.

The Government reserves the right to negotiate training requirements due to existing experiences or unanticipated challenges that new technologies may bring.

The Government desires that the Contractor provide annual technology debriefs and roadmaps by Engineers who are acknowledged leaders or experts in their field.

C.8 Project Plan Requirements

C.8.1 Project Management

The Contractor shall be responsible for all functions related to managing the NOAA HPCS R&D system as a single project. NOAA will retain responsibility for project oversight and directing project activities (such as resource allocation) that directly impact the Government. The Contractor shall establish communication avenues to keep NOAA managers apprised of daily status and alerts. A formal project management review process shall also be established as well as described below.

Reference materials for standard Government Information Technology project management are available on the Internet, for example:

<http://cio.doe.gov/ITReform/sqse/publications.htm#Checklists>

The Contractor shall appoint a Project Manager and assistants as required. This management team shall be responsible for meeting NOAA's technical and business requirements (including security), project planning (cost, schedule quality), resource mapping, identifying dependencies and support issues and managing Sub-Contractors. The Project Manager shall hold an Annual Meeting for NOAA management in Silver Spring, MD to discuss project status and plans for the coming year.

C.8.2 Transition to "One NOAA"

Transition to "One NOAA" requires increased integration of NOAA HPC resources over time to provide additional flexibility and robustness. Elements of the long-term vision include:

- 1) the ability of users to access any computational platform in order to:
 - a) achieve code interoperability with good performance

- b) provide continuity of operations
 - c) respond to changing programmatic requirements
 - d) provide similar programming environment and tools across platforms to minimize user re-training.
- 2) the ability of users to access data from any platform of the HPCS and to migrate data from any existing archive to future archives
 - 3) an ongoing collaboration with the Contractor to produce additional business processes that promote integration

Features of the NOAA R&D HPCS that promote this vision may include a single security profile, including single sign-on, and a single user interface into all HPC resources, which may include a single batch scheduler or a metascheduler. NOAA recognizes that not all elements of the vision may be available on the system provided during the term of the base contract, or that some elements may compromise performance unacceptably. In that case, NOAA may choose to forego or delay implementation of one or more of the elements of this vision.

NOAA considers the risk of transition to the full suite of these capabilities at day one of the contract to be very high due to both the immaturity of the underlying technology for distributed operations and the significant business process changes required of users and institutions.

C.8.3 Documentation

The Contractor shall provide adequate documentation for maintenance of all Project Management and the Configuration Management system. This includes, but is not limited to, baseline configuration, implementation methodology, a formal project plan, and documentation of administrative tasks.

C.8.4 Configuration and Change Management Plan

The Contractor shall establish a Configuration Management (CM) process including a Configuration Management Control Board (CMCB) [that includes both Contractor and NOAA Government staff. The established board must be approved by the Government prior to the initial board meeting.](#) The Government will co-chair the CMCB. The Government desires that the CMCB and CM processes integrate into existing Government Control Boards and processes where applicable. CM plans (at the minimum for network, hardware and software (see Section C.4.2)) shall be developed by the Contractor and presented to the CMCB for approval. Not less than monthly CMCB meetings shall be held to review system status and plans for the future. The Contractor shall be responsible for preparing an Annual Configuration Management Plan for presentation to NOAA management at the annual meeting in Silver Spring, MD. The CMCB will be responsible for oversight of benchmark workstreams, modifications to the benchmarks and for scheduling system time in support of benchmarking. The Government will retain final approval authority for all CMCB actions.

C.8.5 Transition Requirements

C.8.5.1 The Contractor shall develop and submit, within 60 days of a written request by the Government, a comprehensive Transition Baseline Report. The goal of the Transition Baseline Report is to avoid disruption of the day-to-day conduct of HPC computing while achieving a smooth and orderly transfer of responsibility from the Contractor to a successor Contractor. At a minimum, the Transition Baseline Report shall contain all information related to current system configuration, performance, support requirements, issues, risks, and other information deemed pertinent by the Government. This plan shall be submitted to the Contracting Officer and COTR.

C.8.5.2 Thirty (30) calendar days prior to the completion of this contract, an observation period shall occur, during which time personnel of the successor Contractor may observe operations and performance methods of the outgoing Contractor. This will allow for orderly turnover of facilities, equipment, and records and will help to ensure continuity of service. The outgoing Contractor shall not defer any requirements for the purpose of avoiding responsibility for or of transferring such responsibility to the succeeding Contractor. The outgoing Contractor shall fully cooperate with the succeeding Contractor and the Government so as not to interfere with a smooth transition and ongoing accomplishment of the Government's work.

C.8.5.3 Thirty (30) calendar days prior to a scheduled midlife system upgrade the Contractor shall submit a written plan to the Government describing how the upgrade will be implemented. At a minimum, it shall contain a schedule, information related to configuration, performance, support requirements, issues, risks, and other information deemed pertinent by the Government. This plan shall be submitted to the Contracting Officer and COTR.

C.8.5.4 The Government's data is essential to its mission. Therefore thirty (30) calendar days prior to all data migrations, the Contractor shall submit a comprehensive Transition Plan. At a minimum, it shall contain an overview, a schedule, information related to configuration, performance, support requirements, issues, risks, a cost/benefit analysis and other information deemed pertinent by the Government. This plan shall be submitted to the Contracting Officer and COTR.

C.9 Contract Options

C.9.1 Option Period

The Base Period of the contract is FY2006-FY2009. The Government requires an option to extend the contract for another four years from FY2010-FY2013. Requirements are listed in Section C. The Government requires a draft mid-life proposal to be given to the COTR no later than October 2008. The draft proposal shall include an overview presentation, technology roadmap, minimum "X-factor" and SLT improvements, and shall set initial expectations for minimum hardware deliveries. Proposals are due no later than April 2009.

C.9.2 One-year extension of Base Period

After the four-year Base Period, for FY2010, the Government requires the option to extend the lease for operations, maintenance and the equipment for one additional year in quarterly increments. The Government estimates that the annual funding for this extension to be half of the FY2009 funding level (see Section C.4.3).

C.9.3 One-year extension of Option Period

After the four-year Option Period, for FY2014, the Government requires the option to extend the lease for operations, maintenance and the equipment for one additional year in quarterly increments. The Government estimates that the annual funding for this extension to be half of the FY2013 funding level (see Section C.4.3).

C.9.4 Additional R&D HPCS Augmentations

Over the life of the contract there may arise situations that will require the R&D HPCS to be augmented. One such situation might result if NOAA identifies a new requirement for HPC that did not exist at the time of contract award. A second situation might result if NOAA were to enter into an inter-agency agreement with another federal agency to supply that agency with computational resources. In order to accommodate these, and similar situations, the contract includes, as an option, system components and workstreams to augment the R&D HPCS as necessary. The optional system augmentations are comprised of several major elements to include workstreams (CLINS 0010A through 0010J inclusive), compute systems (CLIN 0010K), storage systems (CLIN 0010L), network (CLIN 0010M and 0010R), software (CLIN 0010N), workstations (CLIN 0010P), and maintenance (CLIN 0010S) as identified in Section B.

If additional workstream(s) should be acquired under this option, it is understood and agreed that the actual configuration acquired may be different than the workstream configurations identified in Section C of the contract. In such case(s), the actual configuration being acquired will be negotiated.

Furthermore, in addition to the items identified in Section B of this contract, the Government may, during the term of this contract, require goods and services not otherwise specifically identified, but within the scope of this contract. Should requirements arise, these additional goods and services will be negotiated under the Changes Clause of this contract.

C.9.5 Engineering Support

The Government requires an option to acquire additional expert-level engineering to address impending needs for a given workstream. Due to off-site meetings or presentations, some travel may be required as needed.

C.9.5.1 Applications Analyst

The Government requires an option to acquire senior applications analysts for porting, tuning, optimizing, and developing scientific applications in numerical weather prediction or climate prediction. The Applications Analyst shall work with

NOAA Application engineers, and Scientists. The Analyst shall have direct access to vendor compiler development teams.

C.9.5.2 Systems/Network/Security Engineer

NOAA may require the following professional engineering services: System, Network, or Security Engineering. Activities may include, but are not limited to the following:

- Install/remove vendor supplied OS and systems software (compilers, batch system, etc) upgrades.
- Manage upgrades for COTS or Community Supported software supplied via the contract.
- Use systems utilities to tune performance.
- Provide utilization and performance information to the extent the systems allow.
- Customize configurations of software and hardware.
- Collaborate with Contractor and Government staff.
- Comprehensively documenting plans and performed activities.
- Diagnose technical problems. Including distinguishing between hardware, software, and user errors.
- Work at an application level in the event job scheduling, interpretive utilities (such as scripting languages), compilers and runtime libraries are suspect.
- Support end users.
- Systems Analysis.
- Evaluate current methodologies and offer expert-level advice on alternatives or target solutions.
- Comprehend and apply security upgrades. Collaborative work with Government IT security staff is essential.

C.9.5.3 Facilities Engineer

NOAA may require professional engineering services in support of its HPC facilities. These services may be architectural, electrical, mechanical or civil engineering services. Activities may include, but are not limited to the following:

- Facility Engineering Consulting Services
- General Engineering Studies
- Specific Engineering Studies
- Design Services
- Design-Build Services

C.9.5.4 Visualization/Data Engineer

NOAA may require the following professional engineering services: Visualization, Graphics, or Data Engineering. Activities may include, but are not limited to the following:

- Generation and manipulation of graphical products.

- Data flow optimization.
- Creating and using tools for accessing, manipulating, viewing and representing the data.
- Data Conversion.
- Metadata manipulation for downstream post-processing.
- Data format evolution and porting.
- End user support.

C.10 Appendix A – Details of current NOAA R&D HPCS

C.10.1 Large-Scale Computing (LSC)

C.10.1.1 Boulder, CO

NOAA's Large Scale Computing located in Boulder consists of five logical clusters based on Intel 2.2GHz Xeon systems interconnected with MyriNet. Each node has two single core sockets and 1 GB of memory. Two clusters contain 62 nodes each, one cluster contains 128 nodes, and two clusters contain 256 nodes. There are four front-ends with 4 GB each. There is also a test bed containing 12 dual-socket single core Opteron systems and 12 dual-socket single core Itanium systems. There is an additional Linux cluster based on 296 dual-socket single core 3.2GHz Intel EM64T nodes most with 2 GB of memory per node, two with 8 GB of memory per node, and four with 16 GB of memory per node. The interconnect is based on Myrinet PCIX-D. This system is approximately equivalent to 85% of the Intel 2.2GHz Xeon systems in terms of total aggregate sustained throughput.

C.10.1.2 Washington, DC

NOAA's Large Scale Computing located in Washington, DC consists of 40 IBM Power4 1.3 GHz Regatta H compute nodes, with a total of 640 dual-core sockets for a total of 1280 processor cores. Each node has 8 Logical Partitions (LPARs) with 4 cores per LPAR. Each core has 1 GB of memory. The interconnect is an IBM SP Switch-2, which provides for multiple plain support for the switching fabric. Each LPAR has 2 SP Switch-2 PCI adapters.

The storage and I/O system is integrated with the switch and contains 4 additional dedicated Regatta H nodes each with 32 cores, configured in 4 LPARs with 8 cores per node. Each core has 2 GB of memory. Each LPAR has two IBM SP Switch-2 adapters.

C.10.1.3 Princeton, NJ

NOAA's Large Scale Computing located in Princeton is designed for the batch processing of computationally intensive jobs. The LSC consists of 11 nodes for a total of 3040 cores. Fast scratch disk is available from each node. The fast scratch disk has eight 1 Gb FC per host and is in a RAID 5+1 configuration.

There are eight SGI Origin 3000s running Irix 6.5.19f and consisting of:

2 hosts with	512 x 600 Mhz MIPS, 512 GB memory, 2.6 TB fast-scratch disk
5 hosts with	256 x 600 Mhz MIPS, 256 GB memory, 0.6 TB fast-scratch disk
1 host with	128 x 600 Mhz MIPS, 128 GB memory, 0.6 TB fast-scratch disk

There are three SGI Altix 3700s running Red Hat Enterprise AS with SGI Propack and consisting of:

- 2 hosts with 256 x 1.5 Ghz Intel, 512 GB memory, 1.4 TB fast-scratch disk
- 1 host with 96 x 1.5 Ghz Intel, 192 GB memory, 2 TB fast-scratch disk

In April 2005, NOAA will receive a performance increment of at least 1.8X above the Origin portion of the Large Scale Computing component.

C.10.2 Post-Processing and Analysis

C.10.2.1 Boulder, CO

There are four visualization servers based on Intel 1.7 GHz Xeon systems. They have dual ethernet connectivity (to the internal cluster network and to the internet via NOAA/FSL's LAN). There are also four cron servers based on 2.2 GHz Xeon systems.

C.10.2.2 Washington, DC

NOAA's HPCS system located at Washington DC contains 16 LPARs that are dedicated to single threaded batch jobs (mostly pre- and post-processing) and for interactive access to the parallel compute environment, which consists of the remainder of this HPCS. Additional post-processing and visualization is performed at several branch servers at NCEP, which are typically Origin 300 or Altix servers.

C.10.2.3 Princeton, NJ

NOAA's Analysis system located at Princeton has two nodes for analysis, post-processing and development work. The Analysis system is comprised of two SGI Origin 3900 systems and one SGI Onyx 3000 with visualization tools.

- 2 hosts with 96 x 600 Mhz MIPS, 96 GB memory, 3.9 TB fast-scratch disk
- 1 host with 4 x 400 Mhz MIPS, 4 GB memory, 0.2 TB scratch disk

The fast scratch disk is comprised of 48 x 1 Gb FC per host that connects to a XFS filesystem in a RAID 5+1 configuration.

The analysis computers are available for interactive use at all times. Parallel applications on up to 16 CPUs may be run for up to 8 hours. For applications that use more than 16 CPUs, use of the Large-Scale Computing is recommended, not required.

C.10.3 Storage and Archiving

C.10.3.1 Home File Systems (HFS)

C.10.3.1.1 Boulder, CO

C.10.3.1.2

There are two Data Direct Networks S2A 6000s. Each has 8 x 1 Gb/s FibreChannel ports. One has 9.2 TB of useable space and the other has 6.7 TB of useable space (this space excludes parity stripes). There is also a Data Direct Networks S2A 8500. It has 8 x 2 Gb/s FiberChannel ports and 50 TB of space.

Files systems are served by a number of Linux-based NFS servers with read/write data rates of 40-50 MB/s each. There are six Dell 2650s with dual 1 GHz Pentium IIIs, one Dell 2550 with dual 933MHz Pentium IIIs, three dual socket single core 2.2GHz Xeon systems, and one dual socket single core 2.6GHz Xeon system. All of the systems have 2 GB memory, one Gigabit Ethernet interface and one 1 Gb/s FiberChannel HBA.

The S2A 8500 is serviced by 8 dual socket single core 2.2GHz Xeon systems with MyriNet, Gigabit ethernet, and one 2 Gb/s FibreChannel HBA. Currently two are serving file systems via NFS, however by September 2005 all will be serving file systems via Lustre.

C.10.3.1.3 Washington, DC

The home file systems on NOAA's HPCS system located at Washington DC consist of 25.6 TB of disk space. Most of this disk space is configured as a General Parallel File System (GPFS), which is integrated with the switch as described in section C.10.1.2 above. GPFS utilizes 2 GB Fibre Channel adapters per storage and I/O LPAR.

C.10.3.1.4 Princeton, NJ

In Princeton, NOAA uses a unified /home filesystem for all scientific computers and workstations. This allows for users to log into various computational resources and have their environment and files follow them. This has required custom .cshrc and .login dotfiles that are designed for the unified home directory.

By default, each scientific user is limited to 10 GB of home directory disk space. /home is used primarily for model source code, batch job printouts, and files for workstation applications. Large data files are kept in /archive (the filesystem that is under HSMS).

The home file system is a 2.3TB SAN from a SGI Origin 3000 server. This file system is being served via CXFS within the HPCS and via NFS to the

scientific workstations. Currently, the filesystem is 60% used by 10 million files. Backups are currently being done via xfsdump to the STK 9940B tapes in the silos.

C.10.3.2 Hierarchical Storage Management System (HSMS)

C.10.3.2.1 Boulder, CO

NOAA's HSMS in Boulder is based upon ADIC's StoreNext software. The robotic component is an ADIC AML/J with 8 LTO tape drives. Managed file systems (disk cache) use the DDN S2A 6000s mentioned above. Currently, 1 TB of space is allocated for two managed file systems. Over 88 TB is managed and is composed of over 4,000,000 files. Access to the HSMS is primarily via locally written get/put scripts, some limited direct access to the managed file systems is allowed.

C.10.3.2.2 Washington, DC

NOAA's HPCS in Gaithersburg, MD has access to two StorageTek Powderhorn silos, each with eight 9940B tape drives and about 1.25 PB capacity. The two silos are fully populated with tape cartridges. NOAA's HPCS in Fairmont, WV has access to one StorageTek Powderhorn silo, with four 9940B tape drives and about 1.25 PB potential capacity. The silo currently holds about 2300 tape cartridges. NCEP's experience indicates, on average, each 9940B cartridge can hold 250 GB. The HSMS is presently considered part of the operational system at NCEP. This storage system consists of tape storage with dedicated disk storage for staging data to be stored to, or retrieved from tape. This system is based on HPSS. Access to this system is through htar and PFTP commands, and a public domain IHSI Unix-like interface. Additional information on HPSS and IHSI can be found at:

<http://www.hpss-collaboration.org>
<http://www.sdsc.edu/Storage/hsi/>

NCEP currently writes about 2TB/day to tape and by the fourth quarter of 2005 all available capacity will have been consumed. The tape archive layout is the same at both the Gaithersburg and Fairmont locations; the storage devices are connected to the computational cluster by dual Gigabit Ethernet connections through a Cisco 6513 router. Each computational cluster supports "interactive" nodes where users control HPSS functions.

C.10.3.2.3 Princeton, NJ

NOAA's Storage Area Network (SAN) in Princeton has 22.5TB of capacity that is accessible through switched 2 Gb/s FC. The LSC in Princeton has eight 2 Gb/s channels per node to this SAN. The AC in Princeton has twenty-four 2 Gb/s channels per node to this SAN.

NOAA's HSMS in Princeton uses SGI's Data Migration software to provide a data archive via the /archive filesystem. This software manages both a disk

and a tape copy of each file in /archive. This is done transparently to the user, under one file name. The HSMS software has a 15.8 TB SAN for file staging.

NOAA's HSM in Princeton is expected to have slightly more than 7PB of data.

To archive a file, a user simply copies or moves the file into their /archive/<user> directory from any node in the HPCS. Once a new file is present in /archive, data migration will automatically make a tape copy of the file. At first, the disk copy of the file is also kept available. If the file is not accessed for several days, data migration will remove the disk copy of the file, keeping only the tape copy. The next time this file is accessed, it will be "staged" from tape to disk. Subsequent accesses will then be from the disk-resident copy of the file. However, files may be moved between subdirectories, renamed, or removed without causing the file to be "staged" from tape to disk.




Figure 2 - Growth in tape-resident files on HSMS in Princeton




Figure 3 - Growth in tape resident storage on HSMS in Princeton

C.10.4 Networking

C.10.4.1 Characteristics of NOAA Data Transports

The table below characterizes the data transport capacity that is available to the contract as part of NOAA's basic network infrastructure.

Table VI

Data transport capacity (without regard to latency)								Production Development
WS 1-3		WS 4-6			WS 7-9			
Data Ingest	Interactive & User data	Operational Data Ingest (Also	Other Data Ingest	Interactive & User data	Data ingest	Interactive & User data	Backup (R/W)	

			applies to OCCS Write)					
GFDL	Yes	Yes	1	1	1	1	1	1
GSFC	Yes	1	Yes	Yes	Yes	Yes	2	YES
Boulder	Yes	1 and 2	Yes	33%, 2	2	Yes	Yes	2

Yes – indicates that the NOAA infrastructure will support the required data transfers.

1 – NOAA has identified dim fiber that may be available for use on Highway 1 behind GFDL and is in discussions with DOE on possible cost-sharing arrangements. The fiber would provide access to the NOAA DC area MAN and should support DWDM or CWDM provisioning at 10 Gbps. Once in place, the necessary bandwidth will be available

2 – NOAA is a participant in the National Lambda Rail, NLR, through the Front Range Gigapop. NOAA is exploring the possibility of provisioning a lambda at 10 Gbps over both the northern and southern routes (for redundancy) between Boulder and the NOAA DC metro MAN via the Mid-Atlantic Terascale Partnership. Once in place, the necessary bandwidth will be available.

33% – NOAA has the capability to support about 33% of the required data transport capacity.

NOAA does not guarantee that the latencies over the above WAN connections will be adequate for interactive use.

C.10.4.2 Wide Area Networks (WANs)

Table VII - Capacities

	Internet	Internet2	Other
Boulder	155 Mb/s	310 Mb/s	NLR is expected Q3FY2005
Washington, DC	622 Mb/s	155 Mb/s	
Princeton	15 Mb/s	100 Mb/s	1Gb/s – Princeton University Sayre Hall (Forrestal campus)

Table VIII – Available Bandwidth (approximates)

		Oct. 2004	Oct. 2005	Oct. 2006
Available Wide-Area Network Capacity (Mb/sec.) (list by type below)				
Abilene / Internet 2	BLDR	177	500 ¹	500
Commercial Internet		80	80	100
National Lambda Rail (NLR) ²				
Microwave to Internet2 (PU allowed)= 50 Mb/s	PRTN	25	20	15
Commercial Internet = 9 Mb/s		3	3	3
Commercial Internet	DC MAN	340	340	340
I2		34	500 ³	500
I2	GSFC	3072 ⁴	3072	3072
Dark Fiber		⁵		

Available Wide-Area Network Capacity:**Boulder, CO:**

(1.) WAN capacity assumes timely completion of National Lambda Rail installation and connectivity

Princeton, NJ:

(1.) Assume Princeton University only permits NOAA to use a total of 50% of its Internet2 bandwidth capacity.

(2.) Assume current NOAA usage of total bandwidth will grow at a rate of 15% per year.

¹ Upgrade to 500 Mb/sec will only occur if needed. If NLR is implemented the I2 numbers could change.

² NLR – 10,000 Mbps and or dedicated Lambdas (wavelengths) may be available via NLR in 2005

³ Upgrade to GigE in Q3 2005, 500 Mb/s is an estimate.

⁴ Breakdown: 1Gbps dedicated connection for general users at GSFC. B) 2Gbps with 1Gbps backup (internal R&D effort by various projects).

⁵ See C.10.4.2.2.2

C.10.4.2.1 Boulder, CO.

NOAA's Boulder facility is connected to the Abilene network via the Front Range GigaPOP at 310Mb/s. A 155Mb/s connection is available to Boulder through a commercial Internet provider. NOAA has obtained access rights to connect to the National Lambda Rail.

Note: If the Contractor chooses to use both BLDR-1 and BLDR-2, any LAN connection between BLDR-1 and BLDR-2 must be made with fiber.

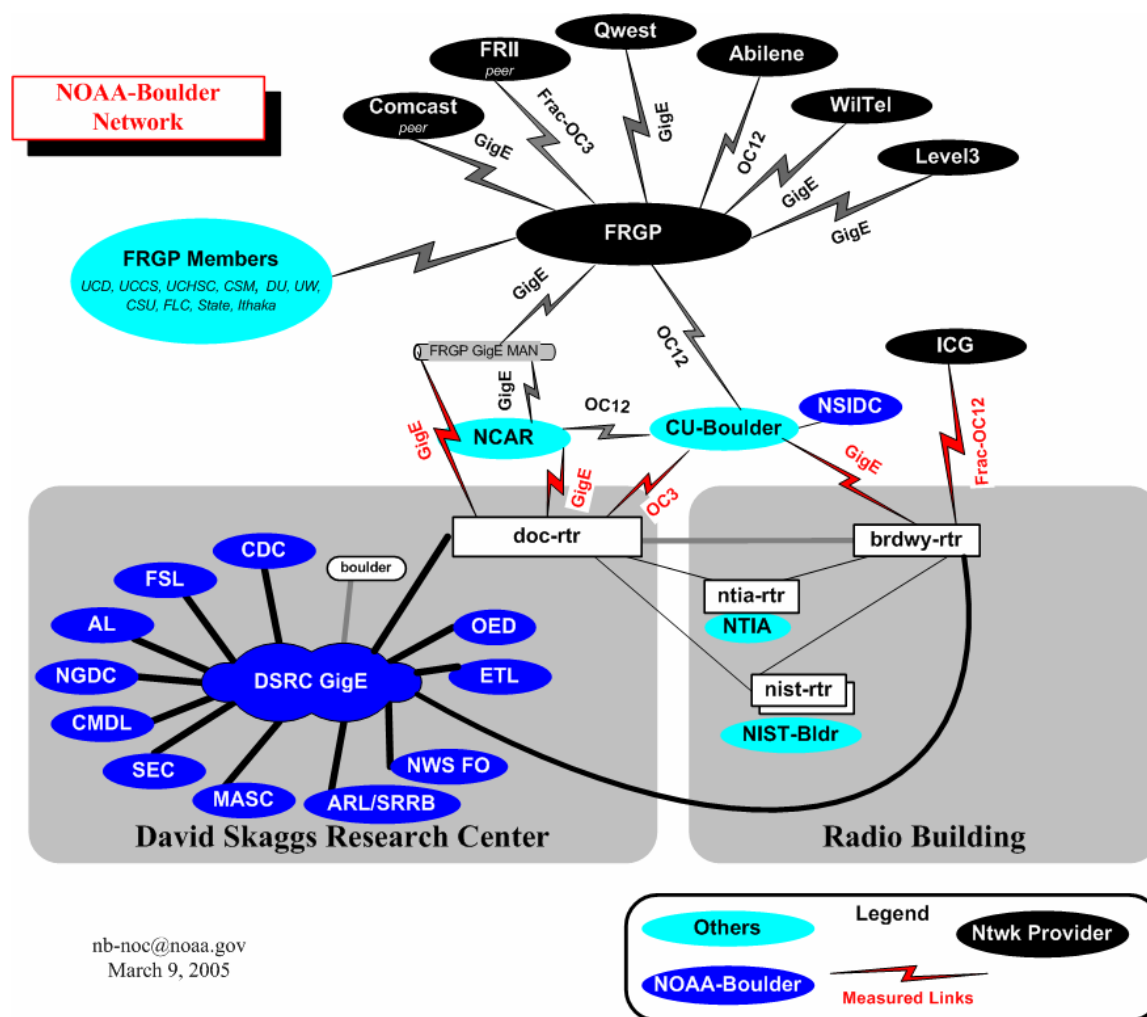


Figure 4

C.10.4.2.2 NOAA's Washington Metropolitan Network (MAN)

NOAA relies on two diverse networks, managed by NOAA's Silver Spring Network Operations Center, to interconnect all Washington Area offices. All offices are connected via Verizon's Transparent Lan Service (TLS), which provides a full duplex connection to a common network "cloud" at speeds of 10Mb/s, or 100Mb/S, or 1000 Mb/S.

In addition locations with very high data transfer requirements are interconnected on dark fiber that NOAA leases from a local provider, Fibergate. All transport on the dark fiber network is currently over gigabit Ethernet. Locations with GigE connections (Silver Spring, Camp Springs, and Suitland), all use that as a primary path and rely on Verizon TLS connections as a back up or failover. The Dark Fiber network also runs to and through the Mid Atlantic Crossroads, Internet2 GigaPop at which point NOAA is also connected to NASA Goddard. Any location on the dark fiber network could be upgraded. Options include:

- 10GigE by upgrading hardware
- Implement a single lambda, by multiplexing.
- Use DWDM and have multiple wavelengths per lambda.

Additional pairs are also available from Fibergate if required.

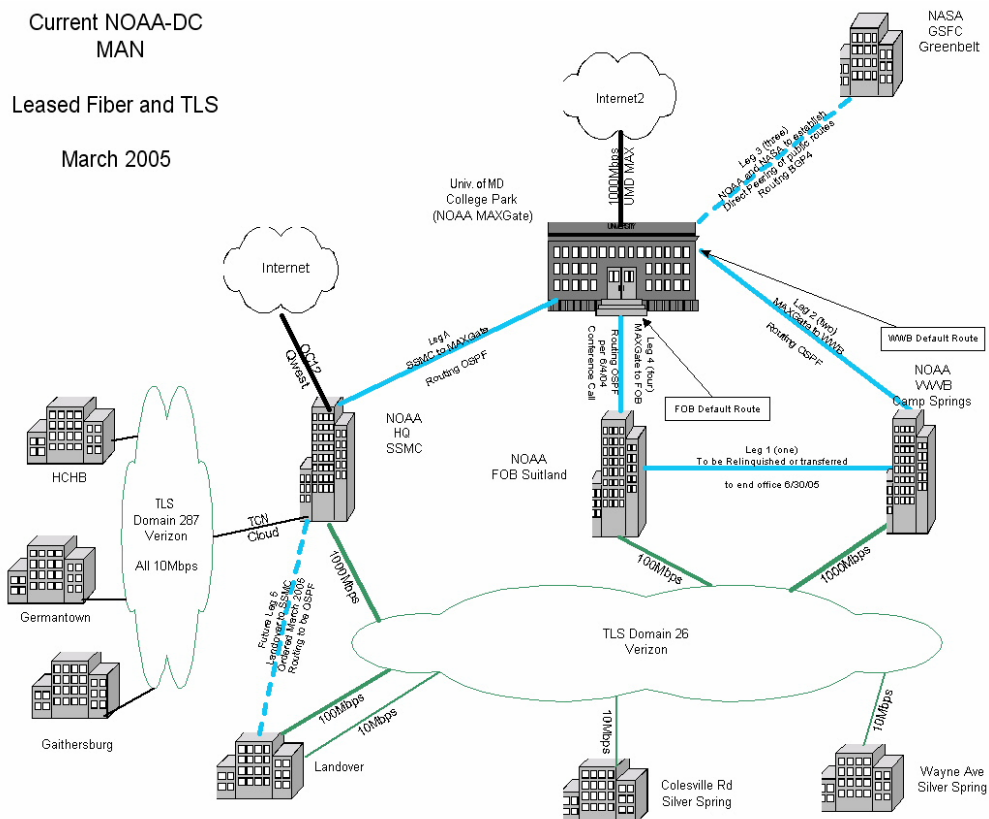


Figure 5

C.10.4.2.2.1 User Connectivity to the Washington MAN

NOAA's Camp Spring, MD facility is connected to Internet2 (dark fiber) with an OC3. An OC12 internet link is presently used close to capacity. This setup is inherited from the previous operational use of the present HPSCS in Washington, DC. It is expected that future HPSCS systems will rely more on Internet2.

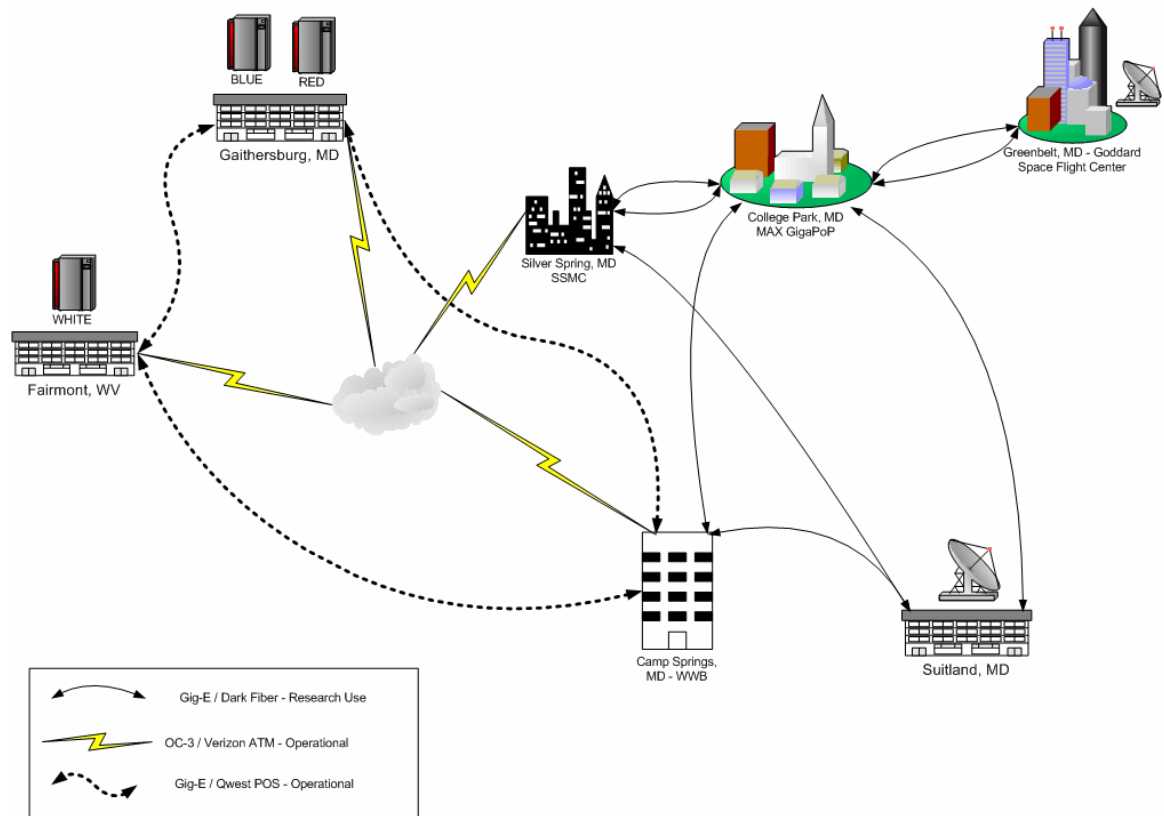


Figure 6

C.10.4.2.2.2 Goddard Space Flight Center (GSFC) connectivity to the Washington MAN

- Building #28 is interconnected to Internet 2 via the Mid Atlantic Crossroads at the UMCP & NASA NISN
- NASA has several pairs of dark fiber between GSFC (Bldg #1) and the MAX. One pair connects GSFC with the MAX Internet2. This is the pair that all the NOAA to/from NASA traffic currently uses
- NOAA has one pair of fiber from the MAX to GSFC (Bldg #1), but it is not being used. This circuit is paid through Feb 06, at which time NOAA will make a determination on whether or not to keep it.
 - Note: Bldg #1 is located approximately 1 mile from Bldg. #28.

C.10.4.2.3 Princeton, NJ

NOAA's Princeton facility is connected to Princeton University's Sayre Hall Building via 1Gb/s Ethernet. NOAA connects to Princeton University through a 100Mb/s microwave link that is accessed through the Sayre Hall Building. The 100 Mb/s Internet2 connection traverses the Princeton University connection. NOAA connects to the Internet through a 15Mb/s connection provided by a commercial Internet provider.

A Cisco 2600 router and a Juniper M7i router are used to connect to the commodity Internet, Princeton University, and subsequently Internet2. These exterior border routers, along with the local DMZ networks, are segmented off from the LAN by redundant Cisco PIX 535 firewalls.




Figure 7

C.10.4.3 MAN to Wide Area Network Connectivity

C.10.4.3.1 Boulder, CO

NOAA Boulder is a joint owner of the Boulder Research and Administrative Network (BRAN). BRAN is an eleven-mile dark fiber metropolitan-area network owned by NOAA, UCAR/NCAR, CU, and the City and County of Boulder. Besides interconnecting BRAN owner sites, it also has key tactical access to commercial communication providers, US West and ICG. NOAA has indisputable access to 24 strands of BRAN fiber. NOAA Boulder and the FRGP have contracted with the National Lambda Rail (NLR) network to attach in fiscal year 2005. Using Dense

Wavelength Division Multiplexing, NLR will provide multiple channels of service at startup including: a 10-Gigabit Ethernet routed circuit to Internet and Internet2, multiple one-Gigabit Ethernet switched circuits to NLR nodes, multiple wavelengths for dedicated point-to-point circuits, and service for the Global Lambda Integration Facility.

C.10.4.3.2 Princeton, NJ

Level3 maintains a dark-fiber run along U.S. Route 1, which is roughly 1/3 mile from the lab across the field. The B-Site of Forrestal Campus is connected to the DOE-supported Princeton Plasma Physics Laboratory (PPPL) through dark fiber. PPPL is on DOE's ESnet wide-area network with a bandwidth capacity of at least OC3.

C.10.4.3.3 Goddard Space Flight Center, Building 28

A dedicated LAN infrastructure (fiber) is available within 25-30 yards to interconnect to 1 to 10Gbps WAN service

C.10.4.4 Local Area Networks (LANs)

All GFE locations support TCP/IP and Ethernet.

C.10.5 Software

C.10.5.1 COTS

C.10.5.1.1 Boulder, CO

NOAA's HPCS in Boulder currently has licenses to use IDL on the visualization servers. The HPCS in Boulder also has licenses for the Portland Group Fortran and C compilers as well as the Intel Fortran and C compilers.

C.10.5.1.2 Washington, DC

NOAA's HPCS system located at Washington DC presently has licenses to use the following COTS software on the LSC: TotalView, Vampir, IDL, Matlab. Also available are the following libraries: Visual Numerics IMSL, ScaLaPack.

C.10.5.1.3 Princeton, NJ

NOAA's HPCS in Princeton currently has licenses to use the following COTS software on the HPCS: CASEVision/WorkShop, MIPSpro f90 Compiler, Etnus TotalView, F-lint, IDL, Legato Networker, Maple, Mathematica, Matlab, MIPSpro f90, NAG Fortran Libraries, NAG Iris Explorer, and S-Plus.

C.10.5.2 Community Supported Software**C.10.5.2.1 Boulder, CO**

NOAA's HPCS in Boulder currently utilizes the following open source software: SUSE Linux, Fedora Linux, Ferret, NCAR graphics, GrADS, and NetCDF utilities.

C.10.5.2.2 Washington, DC

NOAA's HPCS system located at Washington DC presently utilizes the following open source software: GrADS, NCAR graphics, ImageMagick, NetCDF utilities, GEMPAK, and ISMF.

C.10.5.2.3 Princeton, NJ

NOAA's HPCS in Princeton currently utilizes the following open source software: Ferret, Grace, GrADS, NCAR graphics, and NetCDF utilities.

C.10.6 Data Flow Diagrams**C.10.6.1 Boulder, CO**

The following two diagrams show the data flows for the RUC2 backup model and WRF model that are run at the Forecast Systems Laboratory

RUC2 Backup




Figure 8 – RUC2 Backup

WRF


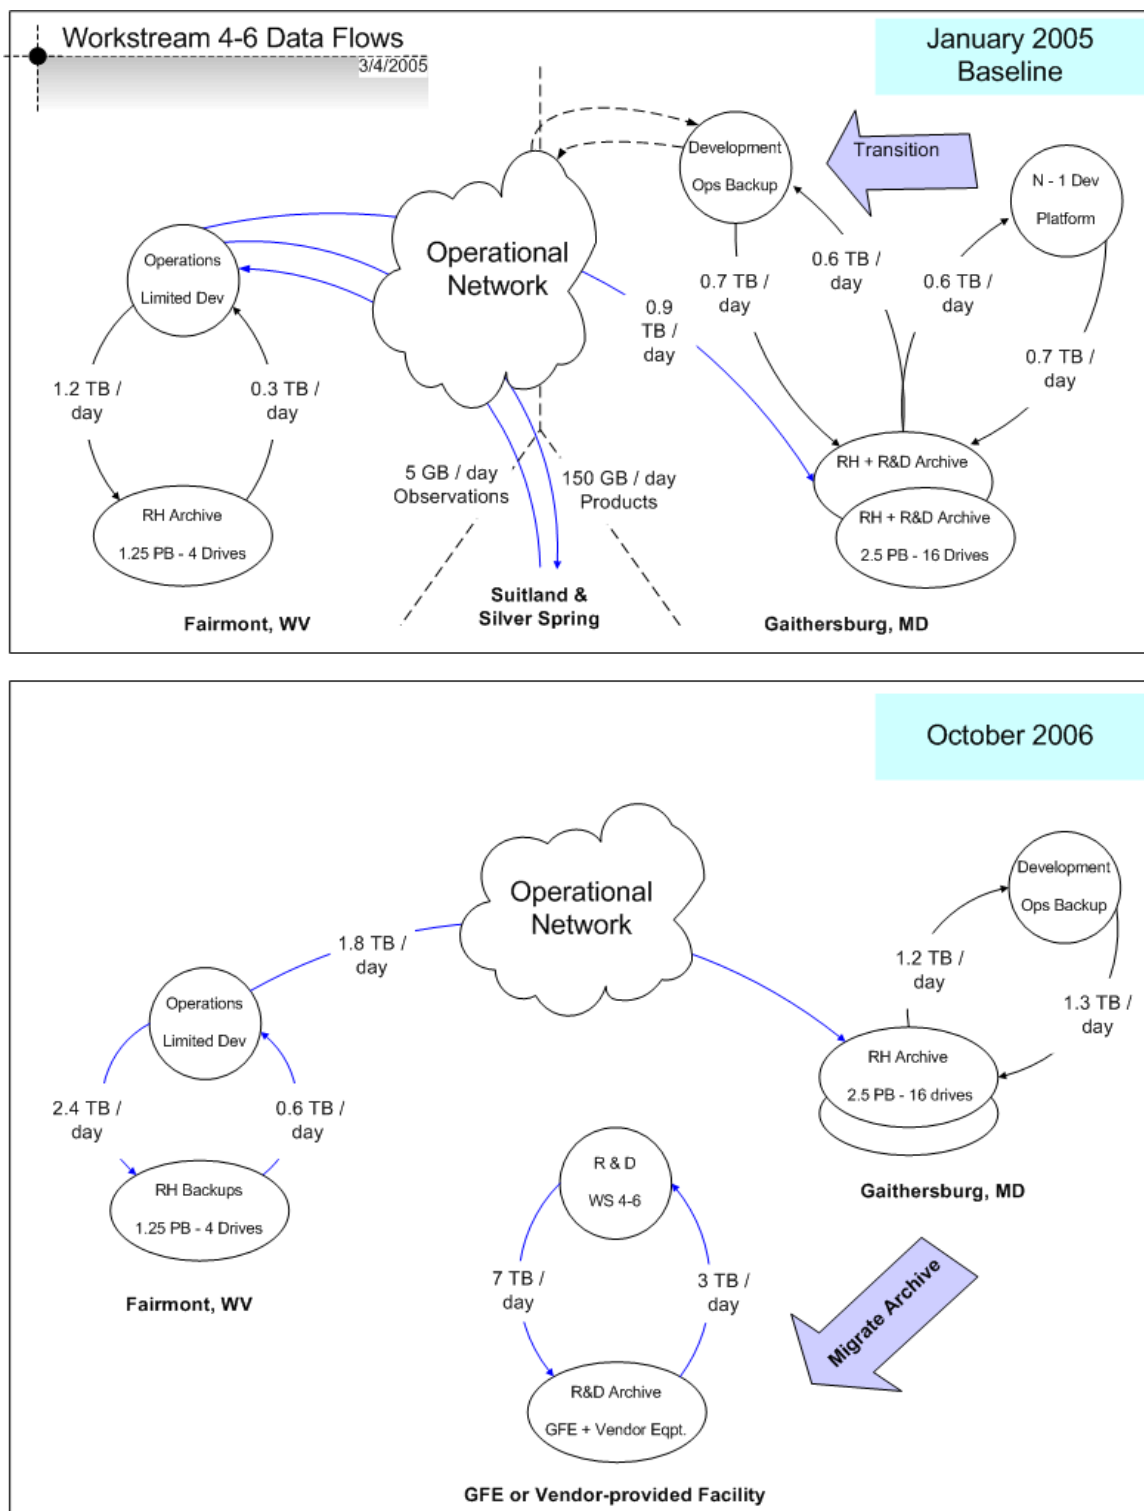


Figure 9 – WRF

C.10.6.2 Washington, DC

The following diagrams show the current data flows and target data flows for workstreams 4-6.



C.10.6.3 Princeton, NJ

The following data flow diagram depicts a single climate run.




Figure 10a – A Representative Climate Run on the Baseline System

C.11 *Appendix B - Available Government Facilities*

This document provides descriptions of four Government-furnished facilities, which are designed to support high-performance computer systems. All four facilities are available for possible use by Contractors to house systems under the NOAA HPC R&D contract, including the availability of facility resources as projected in Section C.5.6 under the assumptions provided therein and subject to the facility terms and conditions provided in Section C.5.6. Two of these facilities, designated as BLDR-1 and BLDR-2, are located in the David Skaggs Research Center (DSRC), 325 Broadway, Boulder, CO. DSRC is owned by the Government and managed by the General Services Administration (GSA). The third, designated as PRTN, is located in Geophysical Fluid Dynamics Laboratory (GFDL) building complex, referred to as the "Princeton Complex" below, at 201 Forrestal Road, Princeton University Forrestal Campus, Princeton, NJ. The Princeton Complex is owned by Princeton University and leased to the Government. The fourth facility, designated as GRBLT, is space located at the NASA Goddard Space Flight Center (GSFC), Building 28, Greenbelt, Maryland 20771. This space will be leased from NASA by NOAA, pending the award decision for this solicitation.

BLDR-1 currently houses NOAA's JET computer system operated by OAR's Forecast Systems Laboratory under a contract with HPTi. The BLDR-1 facility will not be available (with the exception of 80 sqft in October 2005) until October 2006. BLDR-2 is a facility currently under construction at DSRC and is expected to be available for use on February 1, 2006. PRTN currently houses NOAA's HPCS operated by OAR's Geophysical Fluid Dynamics Laboratory under a contract with Raytheon. The GRBLT space is currently used by NASA as general office space.

C.11.1 Layout and Physical Dimensions of Computer Room

C.11.1.1 Computer Room Layout

After Contract award and prior to equipment installation, the Government may negotiate equipment placement in order to address issues such as space utilization and equipment clearances for room access.

Facility Drawings

Contractors requesting facility drawings and related documents that contain sensitive but unclassified information shall be required to have an authorized representative execute a non-disclosure form as required under GSA Order PBS 3490.1. Specific room floor plans are available upon request and are not subject to these regulations.

Figures 11, 11a, 12, 12a, 13, 13a, 14, and 15 referenced below will be provided to Contractors upon request and are not subject to these GSA regulations.

Figure 14 contains sensitive but unclassified information and is subject to the GSA regulations; this figure will be provided to Contractors from whom the Government has received acceptable non-disclosure forms for the facilities shown in the figures (i.e., PRTN).

BLDR-1

Figure 11 shows the computer room layout for BLDR-1. The area within the dotted line will be available in October 2006. This room is only designed for medium-density cooling configurations. Figure 11a shows 80 sq. ft. of space available in October 2005 in support of a possible HSMS.

BLDR-2

Figure 12 shows the planned computer room layout, which is currently in the design phase. It will be designed with both overhead and under floor extreme-density cooling. The uninterruptible power supplies (UPS) will be fed from utility power only (i.e., the room will not be connected to the emergency motor generator). This room will be completed and available for use by the Contractor [on February 1, 2006](#).

Figure 12a shows a revised room layout for BLDR-2 that seeks to address concerns expressed by Contractors in their questions as well as design requirements. The following changes have been made in the design:

- CRAC unit positioning changed (opposing)
- Secondary egress installed
- XDO units moved to accommodate racks that may be up to 48 inches deep
- Full 4 foot “cold aisle” between the fronts of the HPC racks.

Figure 12b shows a revised room layout for BLDR-2 that reveals the current design as of mid-June. The construction contractor RFI process may force some additional small changes. The following changes have been made in the design:

- The Assembly Area room size has been reduced to ~143 sq. ft.
- The two CRAC units on the west side of the room have been removed.
- Electrical distribution on the west side of the room has been removed.
- Various chill water piping chases have been added.
- The west side of the room (within the dotted line ~222 sq. ft.) is available for possible HSM System placement since no additional power or cooling will be offered.
- Offered cooling is reduced to 124 Tons due to Chilled Water Plant limitations, but is still within design specification to cool 350 kVA. Cooling is still commensurate with power.

PRTN

Figure 13 shows the computer room layout for the system as of October 2004, but also showing the silos moved onto the northern hardpan (to be done in

Spring 2005) and room construction on the left hand pan and to the left of the Operators Room that have since been abandoned. References to the front of the room in the following discussion refer to the bottom of the figure (nearest to the Operators Room), while the back of the room is at the top of the figure.

Figure 14 shows the overall Computer Building layout. The rooms adjacent to the Computer Room as shown at the bottom of the figure from left to right are:

- Loading Dock, which is designed to accept deliveries from 18-wheel trucks.
- Storage Room adjacent to the Loading Dock, which also serves as a receiving/staging area for deliveries to the Laboratory.
- Office for Contractor support personnel
- Operators Room, which serves as the control room for computer operations as well as security monitors for the Computer Room as well as outside building access points for the Computer Building and Main Building.
- Printer/User Output Room, which contains local Computer Building printers as well as user output bins.
- Operations Lounge.

GRBLT

Figure 15 shows the space that NOAA will provide in Building 28 at the NASA Goddard Space Flight Center. The figure indicates the planned computer room configuration as the shaded region, which will be converted from two office rooms separated by a hallway.

C.11.1.2 Location of Computer Room and Characteristics of Surrounding Campus

BLDR-1 & BLDR-2

Both BLDR-1 and BLDR-2 are located within the David Skaggs Research Center (DSRC) in Boulder, CO. Due to the large number of visitors, the following website was developed to assist with locality information with regard to the site:

<http://boulder.noaa.gov/>. This website will display maps, local information, driving directions, etc. Regarding specific locations: BLDR-1 is located in the “B” Block of the DSRC, on the 2nd floor, and BLDR-2 is located in the “A” Block of the DSRC, in the Garden level (basement).

PRTN

The Computer Room is in the Computer Building, which is one of two buildings in the PRTN Complex located at 201 Forrestal Road, Forrestal Campus, Princeton, NJ 08540 on Site B of Princeton University's Forrestal Campus. This campus is currently devoted to university research and is intended by Princeton University to be developed in the future as an office park on the U.S. Route 1 corridor. The

nearest airport is the Trenton-Mercer Airport, which is roughly 12 miles away. The nearest highway, U.S. Route 1, is roughly 1/3 mile away.

GRBLT

The computer room is located in Building 28 of the GSFC campus. The GSFC campus is located near Washington, DC.

C.11.1.3 Physical Dimensions of Computer Room**BLDR-1**

Figure 11 indicates the computer room layout for BLDR-1. The entire computer room is 3600 square feet in size, with dimensions of 60 feet by 60 feet. The area within the dotted line in the figure will be available in October 2006 and is estimated to be approximately 2250 sq. ft. Figure 11a shows 80 sq. ft. of space available in October 2005 in support of a possible HSMS.

BLDR-2

Figure 12a indicates the computer room dimensions and layout for BLDR-2. The entire computer room space is approximately 2368 sq. ft. A portion of this space is used for non-HPC equipment and safety egress corridors. 1424 sq. ft. of raised floor space is available for equipment and includes space identified for UPS and CRAC units. An adjacent room contains an additional 156 sq. ft. of raised floor space, and is available as an assembly area.

PRTN

Figure 13 indicates the computer room layout for PRTN. Figure 14 shows the location of the computer room within the Computer Building of the PRTN Complex. The entire computer room is 10,004 square feet in size, with dimensions of 122 feet by 82 feet; this includes the UPS Room, which is located on the hardpan on the right rear corner of the Computer Room.

GRBLT

As Figure 15 indicates, the computer room to be constructed will be taken from two rooms with dimensions of 40'X75' and 40'X50' and a corridor between the rooms with dimensions of 75' x 6'5". The resulting single room will be approximately 5500 square feet.

C.11.1.4 Raised Floor Space**BLDR-1**

A total of 3600 sq. ft. of raised floor is in BLDR-1. 80 sq. ft. of this space will be available in October 2005. A total of 2250 sq. ft. of this space will be available in October 2006.

BLDR-2

A total of 1424 sq. ft. of raised floor will be in available BLDR-2, all of which will be available once construction of the room is completed in late November 2005. See Figure 12a for locations of installed UPS and CRAC unit equipment.

PRTN

The raised floor area in PRTN totals 7052 sq. ft. with dimensions of 86 feet by 82 feet.

GRBLT

All of the provided space will be raised floor space.

C.11.1.5 Non-Raised Floor Space and Equipment Staging Areas**BLDR-1 & BLDR-2**

There is no non-raised floor in either computer room. The computer rooms themselves are used as staging areas.

PRTN

The non-raised floor on the left side of the room is 1476 sq. ft. with dimensions of 18 by 82 feet. Roughly 360 sq. ft. of this non-raised floor, located adjacent to the doors leading to the loading dock, will be available as a staging area for new equipment. The non-raised floor on the right side is 936 square feet, with dimensions of 18 by 52 feet, reflecting reduced non-raised floor space due to the presence of the UPS Room. Figure 13 indicates the planned future location of the five (5) StorageTek silos, four of which are currently located on the raised floor but will be relocated this spring to the locations on the hardpan as indicated in the figure.

GRBLT

There is no non-raised floor space available. The computer room will be used as a staging area.

C.11.1.6 Space for Vendor Personnel, Maintenance Space, and Vendor Storage**BLDR-1 & BLDR-2**

The Government will supply two (2) 150-sq. ft. offices and approximately 100 square feet of lab space for use by the Contractor. None of this space will have furniture or lab benches. The space will have up to four (4) telephones as well as adequate LAN connections.

PRTN

Vendor support personnel (2 Computer Engineers and 2 Software Engineers) currently occupy three (3) rooms in the Computer Building. Equipment storage and maintenance space is currently provided on the left hardpan in the Computer Room and in the UPS Room. Two of the offices will be available to the

Contractor's system support personnel beginning in October 2006. The Contractor may use the non-raised floor space in the southwest corner of the left hard pan to construct additional offices, use as storage space, and/or use for locating equipment that does not require raised floor. Prior to the end of the current Raytheon contract, half of this space (32'x18') will be reserved for Raytheon use. The space to be used for offices that were planned to be constructed (as indicated in Figure 13) on this hard pan space is offered to the Contractor to use to construct offices and/or storage space or else for installing equipment that does not need to be located on raised floors. In addition, the larger Contractor room located on the loading dock side of the Operator Room will be left as is for Contractor use as office space (available to the R&D Contractor after the current Raytheon contract ends) and not divided into two offices as previously planned.

GRBLT

No space other than the space within the computer room is available for offices.

C.11.1.7 Height of Ceiling Above Raised Floor**BLDR-1**

Drop ceiling height = 8.5 feet

BLDR-2

Drop ceiling height = 10 feet

PRTN

Drop ceiling height = 9.5 feet

GRBLT

The drop ceiling height is currently 9 feet above the 18"-high raised floor, but will be 8' 6" if the floor is raised to 24". The final decision on the height of the raised floor will be made post-award.

C.11.1.8 Maximum Allowable Height of Equipment

The following equipment height are mandated by fire codes and cooling limitations.

BLDR-1 & BLDR-2

Due to fire code and cooling limitations, the maximum allowable rack height will be 84". A waiver of this restriction may be allowed, up to 94 inches, depending on the location and placement of the equipment. Modifications to the overhead sprinklers, with approval from the GSA Fire Protection Engineer, may be necessary. It is strongly recommended that vendors who are contemplating the use of tall equipment submit proposed drawings of equipment layouts and specifications to the Government prior to submission of their proposal. The Government will attempt to provide a rough estimate of the costs that will be born by the vendor to accommodate the equipment.

PRTN

Due to fire code regulations, the maximum allowable rack height will be no more than 18" below the drop ceiling height.

GRBLT

Due to fire code and cooling limitations, the maximum allowable rack height is 7 feet. However, special accommodations may be possible, subject to approval of the GSFC Facilities Management Division.

The option of raising the ceiling above the equipment is not permitted because of the presence of obstacles above the current ceiling. Therefore, the following is recommended by FMD as the way to install the equipment, which is assumed in the estimate below to require a 10' by 30' or 300-square-foot area.

The equipment would be installed on the concrete floor with the raised floor tiles and supports removed from the designated area. The concrete floor would need to be sealed and painted. A barrier would be installed around the area to prevent air flow from the surrounding raised floor regions. A railing would be installed on the raised floor surrounding the area, and steps would be installed for access. FMD estimates the cost to the vendor, with multipliers to reach the fully-loaded estimate, to be \$75K for this construction. If the required area is larger than 300 sq. ft., the cost would presumably increase in proportion to the area required.

If the access to the area requires installation of a ramp, rather than steps, then the area required would be increased to accommodate the required ramp with designated slope. A rough estimate of the cost for the ramp is an additional \$18K fully loaded.

The above costs are the Government's best effort to estimate the construction costs for these modifications for the stated area. The vendor will be responsible for the actual costs for these modifications, which will be based on the area required, its location within the room, and any other elements required, such as the possible need for the ramp. As discussed in Section C.5.6.1, after contract

award, the winning vendor will have some latitude in coordinating the plans for this construction with the Government's Phase I site preparations.

C.11.1.9 Space Below Raised Floor

The space below the raised floor is defined as height from the sub-floor base to the top of the raised floor.

BLDR-1

Raised floor height = 12 inches

BLDR-2

Raised floor height = 24 inches

PRTN

Raised floor height = 24 inches (except for the deeper chilled water trench, 4 feet deep, that extends down the middle of the room as indicated in Figure 14).

GRBLT

Raised floor height is currently 18 inches. The plan in Phase I is to raise this to 24 inches, but a final decision regarding this will be made post award.

C.11.2 Installation Characteristics and Raised Floor Capacity

C.11.2.1 Physical Access for Equipment Installation

The Contractor shall provide all equipment needed (e.g., pallet jackets, floor protection, etc.) to support ingress and egress of equipment into the facility during the installation and removal process. Special care must be taken in order to transport the equipment in a manner that does not damage the facility (e.g., scratch the floor, etc.).

BLDR-1 & BLDR-2

Loading Dock - Standard, designed to accommodate 18-wheeled semi-trucks.

Access Path –

Corridors - 8" raised floors

BLDR-1: Doorframes along this path are 83 inches at their low point

BLDR-2: Doorframes along this path are 96 inches at their low point

Elevator – Freight Elevator with 8000-lb. Capacity

BLDR-1: 25' from Loading Dock to the freight elevator and then about 50' to the computer room

BLDR-2: 25' from the Loading Dock to the freight elevator and then about 150' to the computer room, down and up ADA-compliant ramps.

PRTN

Loading Dock – Standard, designed to accommodate 18-wheeler semi-trucks.

Access Path – Equipment passes through an entry door from the Loading Dock into the Storage Room and then onto a hardpan staging area in the south corner

of the Computer Room - a total distance of roughly 30 feet. The two sets of double doors along this path have clearances of 85 inches high by 70 inches wide.

GRBLT

The off-loading of equipment delivered to the site requires a lift-gate truck. Computer room access for equipment will be done via interior circulation corridors.

C.11.2.2 Staging and Assembly Areas**BLDR-1 & BLDR-2**

All staging and equipment assembly shall be performed within the computer room.

PRTN

Equipment may be staged and assembled on the hardpan in the computer room immediately adjacent to the entry from the Storage Room and the Loading Dock. This non-raised floor area is at least 18' wide by 20' deep.

GRBLT

All staging and equipment assembly shall be performed within the computer room.

C.11.2.3 Loading Capacity of Raised Floor**BLDR-1**

Computer Room Raised Floor Specs: ConCore SF 1250 Bolted Stringer:
Concentrated Load: 1250 lbs.; Uniform Load: 300 lbs./ft²; Ultimate Load: 3850 lbs.; Rolling Load: 1000 lbs. (10 Passes)
Raised floor Specs (adjacent hallways): ConCore SF 1250 Cornerlock:
Concentrated Load: 1250 lbs.; Uniform Load: 300 lbs./ft²; Ultimate Load: 3750 lbs.; Rolling Load: 1000 lbs. (10 Passes)

Floor tiles are 2' x 2' and are shown by location in the Figure 11 drawing. The dotted squares represent perforated tiles. Network and power connections are made via 5" round ports within the solid tiles only. All tiles, unencumbered by racks or other equipment, can be relocated for improved airflow, if necessary.

BLDR-2

Computer Room Raised Floor Specs: ConCore SF 1250 Bolted Stringer:
Concentrated Load: 1250 lbs.; Uniform Load: 300 lbs./ft²; Ultimate Load: 3850 lbs.; Rolling Load: 1000 lbs. (10 Passes)

Raised floor Specs (adjacent hallways): Concrete floors

PRTN

Computer Room Raised Floor Specs: Nevamar tile and CEI Bolted Stinger: Static Load: 1250 lbs; Ultimate load: 2900 lbs. Rolling load: 1200 lbs. (10 passes).

Most of the raised floor, indicated as 2' x 2' square tiles in Figure 13, was replaced to the above specs as part of the site preparation for the current Raytheon contract. The exceptions for this raised floor upgrade were for tiles under the following equipment: the two oldest StorageTek silos (when located on the raised floor), the MGE 250-kVA UPS, the Dataflow air-handler along the rear wall of the Computer Room, the Liebert 125-kVA PDU, and the EPE 125-kVA PDU.

Adjacent Non-Raised Floor and Pathway to Loading Dock: Concrete floors.

GRBLT

Computer room raised floor specs: 200 lbs./ square foot. Special accommodations for heavier loads are possible. Contractors shall indicate their needs for raised floor at the GRBLT facility in their facility proposals. A decision regarding the final design of the raised floor will be made post award.

C.11.3 Power Facilities

C.11.3.1 Power Service to Building

BLDR-1 & BLDR-2

480/277-volt main switchboard, rated for 3000 amps and 65,000 AIC. The gear has integrated IMPACC monitoring software. All power is fed from a single utility substation. Power is then distributed via step-down transformers and panel boards.

If the Contractor's power consumption at BLDR-1 and BLDR-2 exceeds the offered amounts in Table C.11.11, then the Contractor shall be required to install electric meters to measure and record power usage of all of the HPC equipment (i.e., excluding facility infrastructure equipment, such as air handlers, that are operated in the Computer Room), both new and GFE, that are operated under this contract. The Contractor shall be responsible for paying all power costs for HPC equipment over the offered amount in Table C.11.11. As indicated in Section H.19, the Government will issue monthly contract payment deductions to provide for this utility payment.

PRTN

The Princeton Complex receives its power from Public Service Electric and Gas (PSE&G) Company. A PSE&G utility substation with a 2500-kVA, 13200-to-4160-volt, oil-filled transformer and 4160-volt switchgear provides electrical service to the Princeton Complex. This substation is located outside of the southwest corner of the Computer Building directly adjacent to the parking lot.

This substation provides power support to the Princeton Complex as well as to several of the other buildings on the B site of the Forrestal Campus. Power usage for these other buildings is primarily for offices, although research activities in some of these other buildings occasionally require substantial power from the substation.

An underground 4160-volt feeder, dedicated to the Princeton Complex, is routed from the campus substation to separate building substations located within the Princeton Main Building and Computer Building. The Main Building substation (1500-kVA 4160-to-480-volt transformer), located on the ground floor of Princeton's Main Building, provides power to the Main Building and to the Chilled Water Plant for the Princeton Complex. A reading from the Main Building substation several months ago showed a peak load of 1000 kW, which probably occurred when we attempted to start-up Chillers #1 and 3 together or operations during the previous Cray/SGI contract.

The Computer Building substation is located in the Transformer Room, the location of which is shown in Figure 14. This substation is comprised of a 4160-volt air interrupter switch, a 1500-kVA silicone-filled transformer, and a 2000-ampere, 480/277-volt main switchboard. The actual rating for this transformer is 1725-kVA, but Princeton University mandates that its usage not exceed 1500-kVA for extended periods of time. The equipment was installed around 1980 when the Computer Building was constructed and provides power to the Computer Building. Currently it also powers the cooling tower for Chiller #4. Princeton is considering installing an A/B switch to provide an option to also run Chiller #4 off this transformer in emergency situations.

On July 22, 2005, the PRTN complex peak power readings were recorded. The PRTN Main Building Transformer had a peak of 752 kW and the PRTN Computer Building Transformer had a peak of 749 kW. The PRTN Complex peak power demand, as recorded at the B-site switchgear, was 1617 kVA and 1464 kW, both of which were reached on July 19, 2005. The rest of the Forrestal Campus served from that switchgear and the 2500 kVA campus substation, had a peak power demand of 869 kVA and 735 kW, both reached July 18, 2005.

The Government is discussing options with Princeton University for immediately increasing PRTN's share of the existing 2500 kVA campus substation to 1600 kW, given that power demand from the rest of the campus seems to be decreasing. In addition, Princeton University and NOAA are negotiating with PSE&G to upgrade the power service to the entire B-site of the Forrestal Campus.

Current discussions involve PSE&G upgrading the underground feeder line connecting the campus to the main PSE&G grid (located on nearby U.S. Route 1) and providing a PMH-9 13,000 Volt switch that would support at least 5000 kVA of power. (PSE&G indicated that this may involve an additional cost to the

customer.) Princeton University would then install a new 5000 kVA substation to service the PRTN complex and the rest of campus. Plans are to upgrade to this configuration within the next 18 months. Such an upgrade would be expected to provide up to 4000 kVA of electric service to NOAA's Princeton Complex. Note, however, that the two currently installed substations within the PRTN complex are 1500 kVA each, effectively capping the PRTN available power at 3000 kVA in the absence of additional infrastructure upgrades.

The Contractor shall be required to install electric meters to measure and record power usage of all of the HPC equipment (i.e., excluding facility infrastructure equipment, such as air handlers, that are operated in the Computer Room), both new and GFE, that are operated under this contract. The Contractor shall be responsible for paying the PSE&G electric utility bill for the Princeton Complex (i.e., power usage associated with both the Main Building and Computer Building transformers). As indicated in Section H.19, the Government will issue contract credits to the Contractor for any power usage above two times the power usage as measured by the above electric meters, where the factor of two includes the power usage associated with cooling the installed system. As indicated in Section C.11.3.2, spreadsheets showing past monthly usage and costs for electric utility bills will be provided to Contractors on request.

The lighting, large mechanical equipment (pumps, A/C units, etc), and some computer equipment are served at 480/277 volts via panel boards located throughout the Computer Building. The building receptacles, small motors, desktop computers, computer room equipment, and similar loads are served at 208/120 volts via step-down transformers and panel boards.

GRBLT

The local utility provider to GSFC is Potomac Electric Power Company (PEPCO). Power to Building 28 originates from the GSFC main substation via two feeders at 13.8KV. The Building 28 north wing service entrance equipment consists of two double-ended metal-clad indoor unit substations, one rated 1500KVA at 480V 3 phase 4 wire, and one rated 1500KVA at 208V 3 phase 4 wire. The available capacity is approximately 1200 KVA on the 480V system and approximately 1200 KVA on the 208V system. Power to the proposed room will consist of one new 800 Amp main distribution panel rated at 480V, 3 phase 3 wire originating from the 480V substation. [The Contractor shall be responsible for the installation of a second 800-Amp feeder line from the basement transformer vault and a second 800-Amp distribution panel if additional power is required.](#) The Contractor shall install a meter to measure the electrical usage of the HPC equipment and report the number of KWHs consumed to NOAA on a monthly basis.

C.11.3.2 Cost of Electrical Utilities Based on Recent Billing

BLDR-1 & BLDR-2

Electrical utility costs for DSRC occupants are currently charged to building tenants according to finished square footage occupied. Based on this algorithm, the electrical cost was \$3.14 per square foot in FY 2004. There is a real possibility of metering electrical usage in the future. In this case, electrical costs would increase substantially. In FY 2004, the average cost for power was \$0.05/kWH.

PRTN

The following table shows usage and expenditure data for the entire Princeton Complex as an indicator of recent electrical utility costs.

Table IX - Total Usage and Expenditures for Electrical Utilities of the Princeton Complex for FY 1999 - FY 2004

	Annual Usage (kWH)	Expenditure	Avg. Cost / kWH
FY 1999	8,826,400	\$713,948	\$0.0810
FY 2000	9,032,000	\$640,544	\$0.0707
FY 2001	8,555,200	\$644,035	\$0.0756
FY 2002	5,976,000	\$431,506	\$0.0722
FY 2003	6,593,600	\$472,405	\$0.0711
FY 2004	7,981,571	\$654,619	\$0.0820

Spreadsheets in MS Excel format containing the Princeton Complex's monthly electric utility cost, power usage (in kWH), and peak demand (in kW) for FY1999 through January 2005 will be provided to Contractors upon request.

GRBLT

Assuming the cost of electricity is 8.1 cents per KWH, beginning in FY07 and with a 3% per annum inflation rate, the Contractor shall include their projected cost for power usage for the computer room in their cost proposal. However, the Contractor will not be responsible for electricity cost for generating chilled water to cool the room.

C.11.3.3 Power Conditioning and UPS Capabilities

BLDR-1

- Power Distribution: Cutler-Hammer Electrical Distribution Equipment (480Volt, 3 Phase) Under floor distribution is accomplished by 50 ft. flexible conduit power whips, fed from wall-mounted breaker panels.
- UPS: 300 kVA Chloride UPS Systems (450 kVA Installed); (250kVA available to R&D HPCS in 2006); 8-Minute Runtime (Full Load)
- Other: Transient Voltage Surge Suppressor (TVSS) Protected
- Emergency Power Off (EPO) Switch Protected

- The DSRC Building standard for UPS Systems is the Chloride Power Protection System.

BLDR-2

- Power Distribution: Cutler-Hammer Electrical Distribution Equipment (480Volt, 3 Phase)
- Under floor power distribution will be accomplished by fixed receptacles attached to rigid conduit and mounted one foot below the floor tiles.
- UPS: 350 kVA UPS System (500 kVA Installed); 16-Minute Runtime (Full Load)
- Other: Transient Voltage Surge Suppressor (TVSS) Protected
- Emergency Power Off (EPO) Switch Protected
- The DSRC Building standard for UPS Systems is the Chloride Power Protection System.

PRTN

Power Distribution Units (PDUs) [shown in dark blue in Fig. 13]:

- Three (3) United Power 225-KVA PDUs (owned by Government/Princeton University), each containing 4 breaker panels, with each panel containing 42 slots
 - Liebert 125-KVA PDU (owned by Government/Princeton University), containing 3 breaker panels, with each panel containing 42 slots
 - EPE 125-KVA PDU (owned by Government/Princeton University), containing 2 breaker panels, each panel containing 42 slots
- UPS [shown in brown in Figure 13]:
- 500-kVA UPS MGE cabinet and battery bank (owned by Raytheon), Model No. '72-130104-00 EPS 6500/44, 66' with Serial No. 69937-01 (Originally installed in 1997)
 - 225-kVA UPS MGE cabinet and battery bank (owned by Government/Princeton University), Model No. '72-130101-01 EPS 6225/44, 66' with Serial No. 200834-01 (Originally installed in 1998)

GRBLT

No power conditioning or UPS capabilities are provided.

C.11.3.4 Backup Power Generator Capabilities**BLDR-1**

- All Equipment fully backed up by emergency backup generators.
- Current Generator is 1250-KW Cummins Diesel. 1800-gallon tank capacity allows 24 hours of operation

BLDR-2

The DSRC EM Generator has reached its rated capacity, and no further load will be allowed. Therefore, the BLDR-2 facility will not have EM generator backup.

The UPS systems will be designed to sustain a 15-minute outage at full load. The recent power loss history at DSRC is presented in Section C.11.7.1 below.

PRTN

Natural-gas-fired, 75-kW generator: activates automatically when power is lost. This generator is only designed to allow pumps within the chilled-water cooling system to continue to run during a brief power outage.

GRBLT

No backup capabilities are available.

C.11.4 Cooling Facilities**C.11.4.1 Cooling Service to Building****BLDR-1 & BLDR-2**

Building operates three (3) 470-ton chilled-water cooling systems. Two of the systems are on EM generator power, and one is fed strictly from utility. There are three cooling towers, each sized to match the capacity of a single chiller. There are three primary pumps, three secondary pumps, and three condenser water pumps. The chiller plant is located in a mechanical room that is adjacent to BLDR-2. Cooling is delivered at 42°F. There is a rough-in plan for an additional chiller and space for an additional cooling tower in the future, but no plans have been authorized to expand the chiller plant at this time. The chiller plant is currently operating at full capacity.

PRTN

Primary Chilled Water Plant, located in the mechanical room and tower bay southeast of the Transformer Room (see Figure 14), contains two chillers [with the following name plate ratings](#):

- 400-ton Carrier centrifugal chiller (designated Chiller #1)
- 350-ton York centrifugal chiller (designated Chiller #3)
- Baltimore Air Coil Cooling Towers

Chiller #1 was installed in the spring of 2000 along with new cooling towers and pumps. Chiller #3 was installed in 1996.

Secondary Chilled Water Plant, located in the Transformer Room (see Figure 14) [with the following name plate rating](#):

- 225-ton Carrier centrifugal chiller (designated Chiller #4)
- Baltimore Air Coil Cooling Tower [\(the cooling tower is actually rated at 255 tons\)](#)

Chiller #4 was installed in 1979. Cooling tower replacement was completed in 2004.

[By the summer of 2006, the Government will install a chilled water loop that will allow the Princeton Complex to operate all three chillers to support its cooling load. This installation will be done using funds outside of the R&D contract. The chilled water piping modifications will utilize a Primary / Secondary pumping](#)

system arrangement that will allow all three existing chillers to operate independently or together in any combination to match the load of the building (so long as sufficient power is available). All three chillers will be interconnected on a 12" primary loop that will be hydraulically independent of the secondary loop. This arrangement will provide constant flow through each chiller, which will allow the maximum heat transfer from each machine.

A new variable speed pump will be provided to supply all of the chilled water distribution from the primary loop throughout the facility including the building comfort conditioning, the computer room ventilation air and the computer room water cooled stand-alone computer air conditioning equipment. All three-way chilled water valves within existing terminal equipment will be required to be replaced with two-way valves and a differential pressure gauge will be required to be located at the farthest point from the chilled water plant to monitor the differential pressure between the supply and return piping. The differential pressure gauge will be calibrated to vary the speed (and hence flow) of the secondary pump to match the chilled water demand of all equipment.

Three constant-speed back-up pumps shall be provided as part of the secondary pumping distribution system to provide redundant chilled water distribution to the computer room stand-alone equipment only in the event of a failure of the variable speed pump.

In addition, the Government intends to replace the current Chiller #4 with a new 250-ton "screw" chiller as part of the chilled water loop installation, using funds outside of the R&D contract. (Note that the other chillers at PRTN are "centrifugal".)

Cooling is delivered to the Computer Room through a six-inch and eight-inch piping system from the mechanical rooms at a temperature of 42° F, plus or minus 2 degrees. The pipe enters the computer room in a trench that is 4 feet deep under the raised floor in the center of the computer room, as indicated in Figure 14. The piping to the Computer Building currently limits the flow of chilled water to the Computer Room and will require redesign if additional cooling capacity is required.

GRBLT

NASA provides campus chilled water service to Building 28 through a 12-inch pipe from its chilled water plant, which is located at Building 24. The chilled water capacity available to the provided computer room is 400-500 gallons per minute. No backup generators are available for the chilled water.

C.11.4.2 Room Air Conditioning Capabilities

BLDR-1

- 90-Ton Liebert Downdraft Computer Room Air Conditioners (CRACs) (De-rated for altitude) (120 Tons Installed via 4 x 30 Ton units)
- Cooling powered from Emergency circuits.

BLDR-2

- 90-Ton Liebert Downdraft CRAC (De-rated for altitude) (120 Tons Installed via 4 x 30 Ton units)
- Seventeen (17) 4-Ton Liebert Extreme Density Overhead (XDO16) Systems
- Cooling powered from UPS

PRTN

- Seven (7) 35-ton Dataflow/APC CRACs (owned by the Government/ Princeton University)
- Other smaller air handling units - one located in the ceiling of the Printer Room and two in the UPS room

The primary source of humidification for the Computer Room is an electric steam boiler, located in the building tower on the northeast corner of the Computer Building.

GRBLT

Six 33-ton Liebert air handlers will be provided as site-restricted GFE equipment. See Section C.12.

C.11.4.3 Availability of Plumbing for Chilled Water**BLDR-1 & BLDR-2**

All chilled water plumbing is installed to accommodate the Computer Room Air Conditioner (CRAC) units and the Extreme Density Overhead (XDO) systems. No additional plumbing capability or capacity is available.

PRTN

The computer room has operated several water-cooled computer systems (most recently - Cray T932, T94, and T3E) and thus has been previously configured to support water-cooling, including a chilled-water trench running down the middle of the raised-floor area of the computer room, as indicated in Figure 14.

GRBLT

Plumbing will be provided to the room for use with air handlers or to cool water-cooled computer systems.

C.11.5 RESERVED**C.11.6 Fire Alarm and Fire Suppression Capabilities****C.11.6.1 Fire Alarm and Suppression Systems****BLDR-1**

Computer rooms has FM-200 Clean Agent fire suppression with VESDA fire detection systems. The room has wet-pipe overhead sprinklers with 155°F trip point. CO₂ bottles are located in the room with FE-36 bottles surrounding the room. Alarms are tied to the building alarm system.

BLDR-2

Computer room will have Clean Agent fire suppression with high sensitive fire detection system. Room has wet-pipe overhead sprinklers with 200 degree F trip point. CO₂ bottles are located in the room with FE-36 bottles surrounding the room. Alarms are tied to the building alarm system.

PRTN

The PRTN computer room currently is protected by a high-voltage fire alarm system and a wet-pipe water sprinkler system. However, this coming spring/summer, these systems will be replaced with the following: (1) a low-voltage alarm system with advanced alarm panel in the room adjacent to the computer room, (2) a dry-pipe water sprinkler system, (3) plumbing and other infrastructure to support future use of a gaseous fire suppression system (such as FM-200), and (4) a VESDA early smoke detection system.

GRBLT

Sprinklers are installed in the computer room. Smoke detectors are installed above the drop ceiling and under the raised floor.

C.11.6.2 Fire Extinguishers**BLDR-1 & BLDR-2**

Fire extinguishers located within the computer rooms are CO₂ bottles. Fire extinguishers located within the hallways and computer room access areas are FE-36 gas bottles.

PRTN

Fire extinguishers are located at various locations throughout the Computer Room and in the Storage Room adjacent to the Loading Dock. Following the NCEP fire event in 2000, PRTN reviewed its fire procedures and eliminated any dry-chemical fire extinguishers from the entry paths to the Computer Room. The lab also met with the supporting fire departments, both Plainsboro and Princeton Plasma Physics Laboratory (PPPL) fire units, to explain the dangers of dry chemical extinguishers to computer equipment and to verify fire suppression procedures to be followed during fire emergencies.

GRBLT

Fire extinguishers are available, but they are not suitable for computer room fires. The proper fire extinguishers will be provided by the Government as part of the Phase I build-out.

C.11.6.3 Response Time of Local Fire Department

Response time is defined as the time duration between fire alarm activation and the arrival of fire department personnel and equipment at the site.

BLDR-1 & BLDR-2

Boulder Fire Department provides fire-fighting capability. A station is located approximately one mile away and the response time is between 4 to 6 minutes.

PRTN

The Princeton Complex has measured the response time of the fire department (Plainsboro Fire Department as first responders, backed up by Princeton Plasma Physics Lab) to average 10 minutes (minimum of 7 minutes and maximum of 14 minutes), based on 6 events (2002-2004). These events were caused by unrehearsed false alarm events that were caused by faulty alarm sensors. [The Princeton Complex will be replacing the entire fire alarm system during the winter of 2004-05.]

GRBLT

The Prince George's County Fire/EMS Department can be on-site in under six minutes. Goddard security has a 2-minute response time.

C.11.7 Special Capabilities**C.11.7.1 Capabilities for Continuous Operation**

The following describes any capability that the facility may have to permit continued operation during a power outage or to ride through power anomalies.

BLDR-1

All electrical systems are fully backed up for uninterrupted service to cover for any length of power outage. The CRAC units are powered from emergency circuits, and will experience a shutdown until the EM generator comes online and transfers power (usually one to two minutes.) at which time they will automatically and sequentially restart. Once the EM generator starts, it will continue to operate until clean and stable utility power is maintained for 30 minutes, at which time the generator will transfer the load back to utility power and shut down.

BLDR-2

No EM generator capacity is available for this facility; therefore, all equipment will be placed on to UPS power fed from utility. Should a lengthy outage occur, the UPS would provide uninterrupted service, until which time a graceful shutdown of computer equipment is required. CRAC units and cooling systems will remain on

until all equipment is off and all static air has been forced from the compute nodes. The CRAC units will then shutdown, prior to exhausting the battery capability of the UPS systems.

The DSRC outage history is in the following table and shows one outage per year that would mandate a shutdown of the computer room.

Table X- Recent History of Power Loss Episodes at DSRC

FY2003		
DATE	DURATION	Resulting Downtime
26 Mar 03	13 Minutes	0 Min
13 May 03	13 Minutes	0 Min
8 Sep 03	30 Seconds	0 Min
11 Sep 03	2 Minutes	0 Min
FY2004		
10 Jun 04	52 Minutes	0 Min

PRTN

PRTN's facility environment is designed primarily for riding through a brief power outage, assuming the outage is sufficiently brief that the chillers will be able to automatically restart so that the temperature and humidity conditions in the computer room are maintained within acceptable limits.

The following table indicates environment related outages for the last three fiscal years. A review of operator logs for these years showed 29 documented power fluctuation episodes (14 in FY02, 11 in FY03, and 4 in FY04) in which the UPSs and PDUs maintained conditioned power to the computer systems and the chillers either continued to operate or else shutdown and then restarted automatically.

Table XI - Recent History of Environment-Related Outages for PRTN

FY2002	
23 Jul 02	Mechanical (chiller) outage of 53-minute duration with 36 minutes of system recovery time
FY2003	
25 Feb 03	Power outage of 103-minute duration with 224 minutes of system recovery time
FY2004	
	No outages

GRBLT

No capabilities for continuous operation are available, although outages are not common.

C.11.7.2 Duration of Ride-Through at Full Load**BLDR-1**

Not Applicable – Maintain operation during power outage

BLDR-2

Designed for graceful shutdown during a power outage lasting longer than 10 minutes, with CRAC units shutting down last.

PRTN

All of the batteries in GFDL's two UPSs, the 500-kVA UPS and the 225 kVA-UPS, are less than one year old. The UPS capacities are both at or near maximum (80%) capacity with the currently installed system and ride-through times are 15-minutes. (Note: The PDUs are also at full capacity.) The room air conditioning units and the chillers do not have generator or UPS support. If the power interruption is brief, Chillers #1 and 3 are likely to restart automatically. This recycling capability, combined with generator backup for the chilled water pumps, permits the facility to ride through brief power interruptions without the system going down. (The government is considering undertaking a chiller plant upgrade which may include replacing Chiller #4 with a unit that has a variable speed drive and has automatic restart capability as well.)

GRBLT

None.

C.11.7.3 Capabilities for Remote Facility Management

The following describes any capability that allows the facility to be monitored and managed remotely.

BLDR-1 & BLDR-2

BLDR-1 has SCADA (Supervisory Control and Data Acquisition) coverage for environmental monitoring. BLDR-2 is projected to have SCADA coverage as well. CRAC units also have building automation tie-in for alarms. Both systems have dial-out capability.

PRTN

Currently the computer room does not have [complete](#) remote monitoring capabilities. Some such capabilities are being [implemented in the Chilled Water Upgrade project, and others are being](#) investigated for possible future implementation.

GRBLT

METASYS equipped environment controls are currently available.

C.11.7.4 Capabilities for Lights-Out Operation

The following describes any capability that allows the facility to operate during off-hours without operators present.

BLDR-1 & BLDR-2

BLDR-1 currently has no operator coverage from 7pm to 7am every day. The SCADA environmental monitoring system will trigger an automated Emergency Power Off (EPO) if the temperature reaches 94F.

PRTN

Currently the Computer Room does not have [complete](#) remote monitoring capabilities. [Some](#) such capabilities are being [implemented in the Chilled Water Upgrade project, and others are being](#) investigated for possible future implementation.

GRBLT

The computer room will be operated remotely.

C.11.8 Physical Security**C.11.8.1 Campus Security****BLDR-1 & BLDR-2**

Both facilities are located on federal property with federal police providing access to the grounds round the clock.

PRTN

Security personnel patrol the Forrestal Campus on a regular basis.

GRBLT

The site is on federal property with access controls.

C.11.8.2 Computer Room Security**BLDR-1 & BLDR-2**

Computer room has restricted access provided by proximity keypads and other controls.

PRTN

Computer Room access is controlled by a cipher lock system and other controls.

GRBLT

Cardkey and brass key will be installed for the computer room.

C.11.8.3 Procedures for Contract Personnel to Access a Government Facility**BLDR-1 & BLDR-2 & PRTN**

All contract personnel shall submit to a background check and will be issued an "affiliate" ID badge. Authorized NOAA and Contractor personnel may escort other NOAA and Contractor personnel as long as the visitors are US citizens. Special arrangements shall be made weeks in advance to accommodate escorted access by foreign nationals.

GRBLT

Access to the NOAA room is authorized by NOAA but controlled by NASA GSFC Facility Security personnel and systems. Authorized NOAA and Contractor personnel may escort other NOAA and Contractor personnel into the NOAA area as long as they are US citizens. Special arrangements shall be made weeks in advance to accommodate escorted access by foreign nationals.

**C.11.9 Estimated Annual Facility Cost of Operation of Computer Room
(Excluding Power Costs)****BLDR-1 & BLDR-2**

Annual costs for rent, maintenance, heating, and other general services were \$42.25/sq. ft. in FY 2004 as paid to the building landlord, GSA.

PRTN

Annual facility costs, excluding electrical utility costs, are estimated to be \$10/sq. ft. This includes: rent paid to Princeton under the Government triple-net lease; facilities services (physical plant maintenance, equipment service, janitorial services, etc.); utilities (other than electrical power); and wide-area network fees.

GRBLT

In addition to the cost of utilities, discussed in Section C.11.3.2, the Contractors shall include the per annum cost of the computer-conditioned space at the GRBLT site which is \$130K per year, beginning in FY07.

C.11.10 Availability of Blueprints**BLDR-1 & BLDR-2 & PRTN & GRBLT**

Blueprints of the facility are available and will be provided to Contractors, subject to their approval of conditions for their use of the drawings.

C.11.11 Projected Availability of Floor Space, Power, and Cooling

The following amounts show projected total resources available at the indicated dates under the assumptions provided. The October 2004 availability is shown for reference purposes only, because the Contractor cannot utilize the BLDR-1 resources until October 2005. The BLDR-2 resources will not be available until [February 1, 2006](#).

Table XII – TOTAL RESOURCES AVAILABLE

		Oct. 2004	Oct. 2005	Oct. 2006
Available Raised Floor Space (sq. ft.)	BLDR-1	0	80	2,250
	BLDR-2	0	1,424	1,424
	PRTN	900	2,700	6,500
	GRBLT	0	0	5,500
	TOTAL	900	4,204	15,674
Available Non-Raised Floor Space (sq.ft.)	BLDR	0	0	0
	PRTN	360	900	2,700
	GRBLT	0	0	0
	TOTAL	360	900	2,700
Available Power for New Hardware (kVA)	BLDR-1	0	10	250
	BLDR-2	0	350	350
	PRTN	250	70	1190
	GRBLT	0	0	900
	TOTAL	250	430	2,690

		Oct. 2004	Oct. 2005	Oct. 2006
Available Cooling for New Hardware (tons)				
	BLDR-1	0	3	72
	BLDR-2	0	124	124
	PRTN	113	51	344
	GRBLT	0	0	260
	TOTAL	113	178	800

Assumptions for Projections

Available Raised Floor Space:

Floor space offered but not utilized by the Contractor's equipment may be used by the Government for other projects, after discussions with the incumbent.

BLDR:

- (1) October 2005 estimate of 1424 sq. ft. is solely in BLDR-2 and assumes completion of computer room design, construction and testing for availability late November 2005.
- (2) 80 sq. ft. of BLDR-1 is available in [February 2006](#) to support an HSMS. (See Figure 11a)
- (3) The increase of 2170 sq. ft. in the October 2006 total reflects additional available space in BLDR-1 that is currently occupied by systems purchased under current HPTi contract. In reality, the new Contractor will be required to coordinate new equipment installation with current system removal so as to minimize loss of compute cycles to the Government.

PRTN:

- (1) Oct. 2004 estimate reflects that the StorageTek silos are still located on the raised floor.
- (2) Oct. 2005 number is an estimate derived from the unoccupied raised floor space for the room layout as indicated in Figure 13a, which shows the assumed positions of systems and Storage Tek silos after the mid-contract installation is complete. (The actual square footage of free space is obviously subject to interpretation to account for obstructions such as room columns and ramps, size of equipment to be installed, etc.)

(3) Oct. 2006 estimate reflects available space occupied by systems leased under current Raytheon contract, because this leased equipment will be removed at the end of the current contract. In reality, the new Contractor shall be required to coordinate new equipment installation with current system removal so as to minimize loss of compute cycles to the Government.

Available Non-Raised Floor Space:

PRTN:

(1.) Oct. 2005 estimate assumes a common staging area of 360 sq. ft. on the non-raised floor adjacent to the exit to the loading dock plus an area of 18'x30' on the left hard pan that would be available to the R&D Contractor for office space and storage. The remaining 18'x32' of non-raised floor space is reserved for storage and work area until the end of the current Raytheon contract.

(2.) Oct. 2006 estimate assumes both the left and right hard pan areas of the Computer Room are available to the R&D Contractor, with the square footage shown including the UPS Room as well. As with the raised floor space, the timing for availability of non-raised floor space requires coordination with the current Contractor to assure smooth transition from the previous contract to the new R&D contract.

Available Power:

Power offered but not utilized by the Contractor's equipment may be used by the Government for other projects, after discussions with the incumbent.

BLDR:

(1.) The projected available power is dictated by the cooling capacity that is available. Currently, the chilled water plant has reached its rated capacity and no further load can be supported without the addition of a new chiller.

(2.) 10KVA of power is available in BLDR-1 in October 2005 in support of a possible HSMS.

PRTN:

(1.) The projected available power for HPC equipment assumes the 2500-kVA campus substation will be upgraded to at least 3000 kVA so the limiting constraint on power to the Computer Room is the 1500-kVA substation servicing the Computer Building.

(2.) Power estimates assume steady state ratings, not book value ratings for systems.

(3.) Oct. 2005 estimate assumes the projected power usage of the current Raytheon system will be 580 kW, which is the maximum (80%) capacity of the two current UPSs.

- (4.) Oct. 2006 estimate assumes the full 1190 kVA is available to the R&D HPC equipment, with other Computer Building power loads assumed to remain the same (e.g., the load of the 7 air handlers was measured at 59 kW). Note that this aggressive power level is due to making available the 290 kVA that was being reserved for the emergency operation of Chiller 4.

GRBLT:

For further information on the total available power, see Section C.11.3.1.

Available Cooling:

Cooling offered but not utilized by the Contractor's equipment may be used by the Government for other projects, after discussions with the incumbent.

BLDR:

- (1.) The chilled water plant has reached its maximum capacity and will lose its N+1 capability on hot days. A failure within the chilled water plant, on a hot day, will force the BLDR-2 computer room to load shed its cooling resources and shut down.
- (2.) 3 tons of cooling is available in BLDR-1 in October 2005 to support a possible HSMS.

PRTN:

- (1.) Assume the building load for the Princeton Complex for the hottest day of the year is 273 tons, based on observation that Chiller #1 alone reached maximum capacity (with the addition of a 10-ton margin of error) on this day when cooling the Princeton Complex and the HPC system that was operating in October 2004.
- (2.) Compute cooling load for the system by multiplying system power usage (converted to refrigeration tons) by a factor of 1.2 to reflect other room heat sources and cooling inefficiencies and assuming a power factor of one for the HPC system.
- (3.) Assume maximum chiller capacity uses Chillers #1 (435 tons) and #4 (250 tons) operating together as the **standard** chiller configuration. Note that during the heat waves of summer 2005, PRTN found that Chiller #1 operates above its name plate rating when using two cooling towers. We understand this is due in part to having this chiller's "impellers" modified. Note also that the chilled water plant will lose its N+1 capability on hot days. A failure within the chilled water plant, on a hot day, may force the PRTN computer room operate in degraded mode.
- (4.) The cooling capacity of the chilled water plant is de-rated by 10% (reduced by 68 tons) from the standard capacity of 685 tons for GFE cooling values. For October 2006 and beyond, the cooling capacity also assumes the installation of

a 250-ton screw chiller to replace the 225-ton centrifugal Chiller #4 so the nominal capacity is increased to 685 tons, which is then de-rated by 10%. Using assumption number 1 above, leaves 344 tons of cooling available.

GRBLT:

For further information on the total available cooling, see Section C.11.4.1. [Figure 11 withheld from Public Website – Available upon request.]

Figure 11 *Projected Computer Room Layout for BLDR-1 in September 2006. The area enclosed within the dashed line indicates the floor space that will be available for rContractor use in October 2006. The small squares shown indicate 2'x2' floor tiles. Vent tiles are indicated by stippled squares.*

[Figure 11a withheld from the Public Website – Available upon Request.]

Figure 11a *Computer Room Layout for BLDR-1 showing 80 sq. ft. of raised floor area available for an HSMS in October 2005.*

[Figure 12 withheld from Public Website – Available upon request.]

Figure 12 *Schematic of Computer Room Layout for BLDR-2 facility that is expected to be available in [February 2006](#).*

[Figure 12a withheld from Public Website – Available upon request]

Figure 12a Schematic of Revised Computer Room Layout for BLDR-2 facility that seeks to address concerns expressed by Contractors in their questions as well as design requirements. See Section C.11.1.1.

“[Figure 12b withheld from the Public Website Available upon Request]

Figure 12b *Schematic of June 2005 Revision of Computer Room Layout for BLDR-2”*

[Figure 13 withheld from Public Website – Available upon request.]

Figure 13 *Computer Room Layout for PRTN as originally planned. Small squares in the figure indicate 2'x2' floor tiles. Vent tiles are indicated by stippled squares. (The construction of the 5 offices on the left hand pan and the two offices to the left of the Operators Room have been cancelled. See Figures 13a and 14 for the actual configuration of office spaces.)*

[Figure 13a withheld from Public Website – Available upon request.]

Figure 13a *Schematic of PRTN Computer Room and adjacent rooms as of October 2005. Equipment in the Computer Room that is not colored is Raytheon equipment that will be removed at the end of the current HPCS contract. The equipment surrounded by red dashed lines, the Verizon telco rack (red) next to the left hand pan and the servers/UPS boxes (light green) and LAN network rack (red) in the lower right, are Government equipment that are not specifically part of this contract. The colored GFE equipment are as follows: air handlers (light blue), PDUs (dark blue), 225-KVA UPS (brown), archival storage (yellow), networking equipment (red), SAN (TP9100) disks (purple), SGI Origin 3800 systems (orange), and SGI Altix 8-core test system (dark*

green). The positions of systems in this figure represent the Government's best estimate at this time, but may change slightly before October 2005.

[Figure 14 withheld from Public Website – Available upon request.]

Figure 14 Computer Building Schematic for Princeton Complex

[Figure 15 withheld from Public Website – Available upon request.]

Figure 15 Schematic showing configuration of planned computer room at GSFC Building 28 as shaded region within the current office layout of the first floor of the building. The location of support columns within the shaded area are indicated by squares.

Request for these figures are to be submitted to the Contracting Officer at the following e-mail address: william.voitk@noaa.gov

C.12 Appendix C – Government Furnished Equipment (GFE)

The following describes Government-owned property that is available to be furnished to the Contractor for the performance of this contract. Because some Government-furnished equipment that is located at the designated sites is only available for use by the Contractor if that facility is used, the GFE lists are separated into two categories: “Site-Constrained GFE”, which is equipment that is only available if used at the designated site; and “Unrestricted GFE”, which is available for use under this solicitation without site restrictions.

Boulder, CO

Site-Constrained GFE

Note: Partial use of this equipment is authorized up to the limits set forth in Appendix B.

Manufacturer	Part/Model Number	Description	Qty	Date Available
Chloride	EDP90/300/250 4x4	Uninterruptible Power Supply UPS Cabinet and Battery Bank (250kVA)	1	10/2006
	EDP70L/100/100 4x4	UPS Cabinet and Battery Bank (100kVA)	2	10/2006
Liebert	FH740C	Computer Room Air Conditioning Chill Water - Down Draft – 30 tons	4	10/2006
Phonetics	SCADA 3000	Environmental Monitoring Sensaphone SCADA 3000	1	10/2006
HSMS				
ADIC	AML/J	Tape Robot	1	10/2005
IBM		LTO-1 Tape drives	8	10/2005
Sun	E450	HSMS Server	1	10/2005
Various		LTO-1 Tape Cartridges	330	10/2005
Data Direct	S2A8500	DDN Raid (~50TB, 150 400GB SATA Drives)	1	10/2006

Unrestricted GFE

Manufacturer	Part/ Model Number	Description	Qty	Date Available
Compute Hardware				
Myricom	M3-E128	Myrinet line card enclosure	15	10/2006
Myricom	M3-M	Management line card	15	10/2006
Myricom	M3-SW16-4DM	Ribbon spine line card	84	10/2006
Myricom	M3-SPINE-8F	Fiber spine line card	44	10/2006
Myricom	M3-SW16-8F	Fiber switch line card	96	10/2006
Myricom	M3F-PCI64B-2	PCI card (NIC)	768	10/2006
Myricom	M3-CLOS-ENCL	14U enclosure for Clos256_256 networks	2	10/2006
Myricom	M3F-PCIXD-2	PCI D card (NIC)	317	10/2006
Myricom	M3-SW32-16F	16 Port line card	20	10/2006
Myricom	M3-THRU-16Q	Line card for 16 quad fiber thru connections	2	10/2006
Myricom	M3-2SW32	Dual SW32 line card	4	10/2006
Myricom	M3-4SW32-16Q	Quad SW32 line card	2	10/2006

Washington, DC

All of the GFE listed in this subsection below will be available October 2006.

Site-Constrained GFE

- None

Unrestricted GFE

Gaithersburg, MD:

Two (2) StorageTek Powderhorn silos configuration (two 9310, 9311, 9330).

Each of the silos includes:

- eight (8) 9940B tape transports
- 5000 9940 tape cartridges

Each silo contains roughly one PetaByte of data.

Fairmont, WV:

One (1) StorageTek Powderhorn silo configuration (9310, 9311, 9330).

The silo includes:

- four (4) 9940B tape transports
- 2300 tape 9940 cartridges

Greenbelt, MD

All of the GFE listed in this subsection below will be available in the summer of 2006.

Site-Constrained GFE

- Six 33-ton Liebert air handlers

Unrestricted GFE

- None

Princeton, NJ

Unless otherwise stated, all of the following GFE for Princeton will be available in October 2006

Site-Constrained GFE

Manufacturer	Part Number	Description	Qty
<i>Computer Room Security Monitoring</i>			
		C-Cure 800	
Compac	4403US	Compac MS 2000 C-Cure 800 OS	1
Emerson	P761VBT0EENC	Monitor C-Cure 800	1
<i>Computer Room Security Monitoring</i>			
		Intellex 16000	
Sensormatic	DVMS DV16000	Intellex 16000 Video recorder system w/CD backup max 16 cameras	1
NEC	FE1250+BK	Monitor Intellex 16000	1
Pelco	CC-3700-S	Color CCD Camera	12
<i>Power Distribution Units</i>			
United Power	PDM4-F3-225-K13-426	PDU 1 (225 kVA)	1
United Power	PDM4-F3-225-K13-426	PDU 2 (225 kVA)	1
United Power	PDM4-F3-225-K13-426	PDU 3 (225 kVA)	1
<i>Uninterruptible Power Supply</i>			
MGE	72-131522-000	UPS Cabinet and Battery Bank 225 KVA	1
<i>Power Distribution Units</i>			
Liebert	PPA125C	PDU Unit (125 KVA)	1
EPE	PD084M48A12-125	PDU Unit (125 KVA)	1
<i>Computer Room Air Conditioners</i>			
DataFlow	CCT-60C4	CRAC 1 (35 Tons)	1
DataFlow	CCT-60C4	CRAC 2 (35 Tons)	1
DataFlow	CCT-60C4	CRAC 3 (35 Tons)	1
DataFlow	CCT-60C4	CRAC 4 (35 Tons)	1

DataFlow	CCT-60C4	CRAC 5 (35 Tons)	1
DataFlow	CCT-60C4	CRAC 7 (35 Tons)	1
APC	CCT-60C4	CRAC 8 (35 Tons)	1
APC	CM-3.0-W-BC-D	CRAC 6 (3 Ton)	1

Unrestricted GFE

Manufacturer	Part Number	Description	Qty
<i>Switch for Shared Storage</i>			
SGI	FC-SWITCH-16	SGI Fibre Channel switch with 16 ports and one power supply	8
SGI	FC-SWITCH-PWR	Optional second power supply for SGI FC switches (8 & 16 port)	8
SGI	FC-SWKIT	Rackmount kit for mounting FC-SWITCH-8 or FC_SWITCH-16 in F1RACK	8
SGI	XSWOPTGBIC	Short Wave Optical GBIC kit containing 6 GBICs	16
SGI	XCOPGBIC	Copper GBIC kit containing 6 GBICs	6
<i>Switch for Tape Drives</i>			
SGI	FC-SWITCH-16	SGI Fibre Channel switch with 16 ports and one power supply	8
SGI	FC-SWITCH-PWR	Optional second power supply for SGI FC switches (8 & 16 port)	8
SGI	FC-SWKIT	Rackmount kit for mounting FC-SWITCH-8 or FC_SWITCH-16 in F1RACK	8
SGI	XSWOPTGBIC	Short Wave Optical GBIC kit containing 6 GBICs	16
SGI	XCOPGBIC	Copper GBIC kit containing 6 GBICs	6
<i>Console and Monitoring</i>			
SGI	SG230-00008	230L Workstation, 667 MHz PIII, 128MB PC 133 SDRAM, 20GB IDE, V3 Gfx 32MB DDR, Red Hat 6.1	2
SGI	91-AB945-001	19" Northern Hemisphere Monitor	2
SGI	91-AD001-001	Keyboard, Mouse, Speakers, Power Cords, Monitor Cable, User Manual	2
SGI	SSU60003	HW on-site support 4hr rspnse, 7x24, years 1-3 for SGI 230 Workstation	2
SGI	SC4-PCP-2.0	Performance Co-Pilot - Performance Monitoring tool for IRIX 5.3 and higher (replaces PCPORIGIN)	1
SGI	SC4-PCPHPC-1.0	Performance C-Pilot Add-On agent for IRIX 6.5 clusters (replaces SC4-PCPARRAY-1.0); requires SC4-PCP-2.0	8
SGI	SV4-PCPCOL-2.0-10	Performance Co-Pilot Collector 10 license pack	1

Hierarchical Storage Mgmt System

SGI	ORIGIN-3800	SGI Origin 3800 server - 64 CPUs (GFE) 600Mhz R14000A cores), 64GB memory, 16 Local Disk Channels (1Gb/sec), 16 Shared Disk Channels(2Gb/sec), 24 Tape I/O Channels, and 2 Gigabit Ethernet Channels. 180GB	1
StorageTek	9310002-0000	POWDERHORN 6000 CART/450 EPH	3
StorageTek	9940L03-0000	9940, Library, Fibre	30
StorageTek	9840L03-0000	9840, Library, Fibre	22
StorageTek	<td>	Titanium, Library, Fibre	5
SGI	FC-SWITCH-16	16 port 1Gb FC switch (8Cu/8Optical)	16
SGI	ORIGIN-3800	SGI Origin 3800 server - 64 CPUs (Lease 1a) 600Mhz R14000A cores), 64GB memory, 16 Local Disk Channels (1Gb/sec), 16 Shared Disk Channels(2Gb/sec), 24 Tape I/O Channels, and 2 Gigabit Ethernet Channels.	1

Storage***HSMS Local Storage***

SGI TP9100 1GB	TP9100 D-Brick with fourteen 36GB 10KRPM Drives	16
STOR-CTRL 128	TP9100A Dual Channel Control Unit with 1Gb FC	16
SGI TP9100 1GB	TP9100 D-Brick with fourteen 18GB 10KRPM Drives	2
STOR-CTRL 128	TP9100A Dual Channel Control Unit with 1Gb FC	2
Empty Rack	I/O Racks with AC Power Distribution	3
FC-SWITCH-16	16 port 1Gb FC switch (8Cu/8Optical)	4

Hierarchical Storage Mgmt System

StorageTek	9310002-0000	POWDERHORN 6000 CART/450 EPH	2
StorageTek	9840	Media Cartridges	6500
StorageTek	9940	Media Cartridges	16500
StorageTek	Titanium	Media Cartridges	5000

Connectivity***Connectivity - Existing***

Cisco	WS-C6509	Catalyst 6509 Chassis	1
Cisco	WS-CAC-1300W	Catalyst 6000 1300W AC Power Supply	1
Cisco	WS-CAC-1300W/2	Catalyst 6000 Second 1300W AC Power Supply	1
Cisco	SFC5K-SUP-5.5.1	Catalyst 6K Supervisor Flash Image, Release 5.5(1)	1
Cisco	WS-X6K-SUP1A-2GE	Catalyst 6000 Supervisor Engine 1A, Enhanced QoS, 2GE	2
Cisco	WS-X6416-GE-MT	Catalyst 56000 16-port Gig-Ethernet mod. MT-RJ	2
Cisco	WS-X6408A-GBIC	Catalyst 6000 8-port GE, Enhanced QoS	2
Cisco	WS-X6348-RJ-45	Catalyst 6000 48-port 10/100, Upgradable to Voice	2
Cisco	WS-G5486	1000BASE-LX/LH "long haul" GBIC	4
Cisco	WS-G5484	1000BASE-SX 'Short Wavelength' GBIC (Multimode only)	36
Cisco	CISCO3640	Cisco 3600 4-slot Modular Router-AC with IP software	2
Cisco	S364C-12.0.4T	IP	2
Cisco	NM-2CT1-CSU	2-Port Channelized T1/ISDN-PRI with CSU Network Module	2
Cisco	NM-1FE1CT1-CSU	1 Port F Ethernet 1 Port Channelized T1/ISDN-PRI with CSU NM	2

<i>Maintenance/Test System</i>			
Manufacturer	Part	Description	Qty
SGI	LS-3700	Altix 3700 Server: base system infrastructure	1
SGI	TP9300-BASE-2C	Fibre Channel controller enclosure, dual 2 Gb FC controllers, 512MB cache, 6 SFPs, 14 drive bays, no drives	1
SGI	TP9500-73GB-E	2 Gbit, 73 GB, 15K rpm LP Disk Drive with Dual FC/AL Ports	14
SGI	DK-TP-001	TP9300 destination kit	2
SGI	TP9300-RMKIT-UNIV	Rackmount kit for OEM racks, mounts to rails 20.5" to 28.5" deep for TP9300 controller enclosure	1
SGI	X-F22-OPT-3M	3 meter 50 micron optical cable w/LC connectors. For use with 2 Gbit switches, HBAs and devices.	2
SGI	LS-3700-4-4	Altix 3700 system: 4:1.5 GHz/6.0 MB, includes 1 IX-brick with 73GB system disk, no memory.	1
SGI	LS37-CBRICK-4-4	Altix 3300/3700 C-brick with 4x1.5 GHz/6.0 MB cores with NUMalink cables, no memory	1
SGI	LS-MEM-8G-166	8 GB standard system memory (8 DDR x 1 GB 166MHz) for Altix 3700 (min. 1 per C –or- M brick required).	2
SGI	73GB-IX BRICK	System disk for use in SGI 3000 series IX-bricks (73 GB/10K SCSI disk) maximum two disks per IX-brick.	1
SGI	SC5-PROPACK-2.4	SGI ProPack v2.4 for Linux (includes base OS-SGI Advanced Linux Environment 2.1.1)	1
SGI	DK-N3P-003	SGI 3000 three-phase power distribution unit for North America. One per SGI 3000 rack.	1

SGI	PCI-FC-1PORT-B	Single port 2 Gbit FC host bus adapter	2
SGI	PCIX-GIGENET-OR	1-port optical Gigabit Ethernet card for Altix	1
SGI	SR5-XVM-IPF-M	Right To Use XVM Standalone Volume Manager on a single Altix 3000 Linux host with up to 16 CPU.	1
SGI	M4-SVCWRNT-1.0	Global Services warranty card, system registration and Services Reference Guide for systems.	1
SGI	M05-CD	CD-ROM Update Media	1

C.13 Appendix D – Terms

C.13.1 Abbreviations

ADIC	Advanced Digital Information Corporation
ANSI	American National Standards Institute
CM	Configuration Management
CMCB	Configuration Management Control Board
CM2	GFDL's Climate Model Version 2
CONOPS	DOC's Concept of Operations acquisition process
COTS	Commercial Off-The-Shelf
CXFS	SGL's shared filesystem for storage area networks (SANs)
DDN	Direct Data Networks
dGB	Decimal Gigabyte (10 ⁹ bytes)
dPB	Decimal Petabyte (10 ¹⁵ bytes)
DOC	US Department of Commerce http://www.commerce.gov
DSRC	David Skaggs Research Center
ENTA	Enterprise Network Target Architecture
FC	Fiber Channel
FISMA	Federal Information Security Management Act of 2002
FSL	NOAA's Forecast Systems Laboratory http://www.fsl.noaa.gov 325 Broadway Boulder, CO 80305
FY20xx	Fiscal Year 20xx
Gb	Gigabit (2 ³⁰ bits)
Gb/s	Gigabit(s) per second
GB	Gigabyte (2 ³⁰ bytes)
GFDL	NOAA's Geophysical Fluid Dynamics Laboratory http://www.gfdl.noaa.gov Princeton University Forrestal Campus 201 Forrestal Road Princeton, NJ 08542
GFE	Government Furnished Equipment
GNU	Unix-compatible software system
GPFS	General Parallel File System
GrADS	Grid Analysis and Display System
GSA	General Services Administration
HBA	Host Bus Adapter
HFS	Home File System
HIMF	Hallberg Isopycnal Model (Fortran)
HPC	High Performance Computing
HPCS	High Performance Computing System (see definition)

HPSS	High Performance Storage System
HIS	Hierarchical Storage Interface
HSM	Hierarchical Storage Management
HSMS	Hierarchical Storage Management System
Htar	HPSS tape archiver
IDL	Interactive Data Language
IEEE	Institute of Electrical and Electronics Engineers
IMSL	International Mathematical and Statistical Library
I/O	Input/Output
IPCC	Intergovernmental Panel on Climate Change
IT	Information Technology
ITS	Information Technology Services
KB	Kilobyte (2^{10} bytes)
KVA	Kilovolt amperes (1000-volt amps).
LAN	Local Area Network
LPAR	Logical Partitioning
LSC	Large-Scale Computing
LTD	Live Test Demonstration
MAN	Metropolitan Area Network
Mb	Megabit (2^{20} bits)
Mb/s	Megabit(s) per second
MB	Megabyte (2^{20} bytes)
MPI-2	Message Passing Interface (latest version)
MTBF	Mean-Time-Between-Failure
NAG	Numerical Algorithms Group
NCAR	National Center for Atmospheric Research
NCEP	NOAA's National Centers for Environmental Prediction http://www.ncep.noaa.gov 5200 Auth Road Camp Spring, MD 20746
NetCDF	Network Common Data Format
NIST	National Institute of Standards and Technology
NLR	National Lambda Rail
NOAA	DOC's National Oceanic and Atmospheric Administration http://www.noaa.gov
NWS	National Weather Service
OAR	Office of Oceanic and Atmospheric Research
OCCS	Operational Central Computer System
OCIO	Office of the Chief Information Officer
OMB	Office of Management and Budget
OPM	Office of Personnel Management
PB	Petabyte (2^{50} bytes)
PDU	power distribution units
PFTP	Program to transfer data from host to host
PM	Preventive Maintenance

POSIX	Portable Operating System Interface
PSE&G	Public Service Electric and Gas
R&D	Research and Development
RAID	Redundant Arrays of Inexpensive Disks
RFI	Request For Information
RFP	Request For Proposal
RUC	Rapid Update Cycle
SAN	Storage Area Network
SGE	Sun Grid Engine
SLT	System Life Throughput
SON	Statement of Need (Section C of this Procurement)
STK	Storage Technology
TB	Terabyte (2 ⁴⁰ bytes)
UDUNITS	Unidata's Units Library
UPS	Uninterruptible power supplies
WAN	Wide Area Network
WRF	Weather and Research Forecast
WRF-CHEM	Weather and Research Forecast with Atmospheric Chemistry
WRF-SI	Weather and Research Forecast with Static Initialization
WS	Workstream

C.13.2 Definitions

Application Memory	The maximum resident set size used by an application process in Mega/GigaBytes.
Availability	The availability level of a computer, component, or device is a percentage figure determined by dividing the operational use time by the difference between wallclock and null time.
Average (Workstream) Throughput Time	The average over all Workstream Instance Throughput Times. See definition of Workstream Instance Throughput Time below.
Communication Fabric	The hardware component(s) supporting MPI message traffic.
Community Supported Software	Open Source Software
computer	computational platform that share a high-speed interconnect (largest useful MPI program)
Core	basic computational unit
Degraded Mode	System operation at less than normal capability due to the loss of hardware or software components on that system.
Downtime	That period of time when all of an HPCS component's workload cannot be accomplished due to a malfunction in the Contractor-maintained HPCS hardware or software, or when the HPCS or a component of the HPCS is released to the Contractor for maintenance. See Section C.6.1.

HPCS	A High Performance Computing System can be one large System or an aggregation of Subsystems. A given System/Subsystem can be further described by individual components. A single component may span multiple Subsystems.
Job Slot	The logical partition of a computational resource supporting an instance of a workstream (e.g. the set of processing elements on which an instance of a workstream is run).
Node	computational platform that shares memory (largest OpenMP or pthreads application)
Null Time	The period of time when the workload cannot be accomplished due to environmental failure at a Government furnished site, such as loss of electric power or cooling, or recovery from environmental failure.
Operational Use Time	The time during which equipment is available to the Government, exclusive of preventative maintenance time, remedial maintenance time, or Contractor-caused machine failure. Partial credit may be given by the Government for equipment operating in degraded mode (for example, when a portion of the sockets, memory, disk, etc. on a computer is unavailable). The Government may declare the entire HPCS down even if parts of the HPCS are available.
Physical Memory per Processor Core	The size of the memory chipset in Mega/GigaBytes supporting a single processor core.
Preventative Maintenance (PM)	That maintenance performed by the Contractor which is designed to keep the equipment in proper operating condition. It is performed on a scheduled basis.
Processor Core	The component of the processor containing the dependently scheduled floating point and integer registers and arithmetic and load/store unit(s).
Processor Socket	The motherboard component designed to receive the processor.
Remedial Maintenance (RM)	That maintenance performed by the Contractor which results from Contractor-supplied equipment or operating software failure. It is performed as required and therefore on an unscheduled basis.
Socket Suite	basic construction unit A set of concurrent instances of a workstream and possibly, a number of given workstreams (see target IT architecture).
Subsystem	The set of components which accomplish a workstream suite. The aggregate of all Subsystems is the HPCS. Synonym for Target IT Architecture.

System Life Throughput (SLT)	A measure of performance and availability delivered for all instances of a given workstream. See Section C.6.1.2 for more information.
Target IT Architecture	The set of hardware and software components which accomplish a workstream suite.
Total (Workstream) Throughput Time	The wallclock time from submission of the first instance of a workstream component to the successful program end of the last instance of a workstream component. In general, multiple workstream types and instances may be running on a particular target IT architecture. The Total Throughput Time for a workstream is defined in the context of all concurrent loads for the target IT architecture.
Workstream Instance	A single instance of end-to-end processing.
Workstream Instance Throughput Time	The wallclock time from submission of the first instance of a workstream component to the successful program end of a given instance of a workstream component. In general, multiple workstream types and instances may be running on a particular target IT architecture. The Workstream Instance Throughput Time for a workstream is defined in the context of all concurrent loads for the target IT architecture.